

# Full Stack Developer

Interview Questions  
and Answers

## Core Concepts

This section focuses on fundamental principles and advanced concepts that an experienced developer should master.

### 1. Explain the event loop in Node.js and how it handles asynchronous operations differently from browser JavaScript.

#### Event Loop Architecture

The Node.js event loop is a single-threaded mechanism that handles asynchronous operations through **phases**. Unlike browser JavaScript which primarily deals with user interactions and rendering, Node.js manages I/O operations, timers, and system calls.

#### Key Phases

- **Timers:** Executes callbacks scheduled by `setTimeout()` and `setInterval()`
- **Pending callbacks:** Executes I/O callbacks deferred to the next loop iteration
- **Poll:** Retrieves new I/O events and executes their callbacks
- **Check:** Executes `setImmediate()` callbacks
- **Close callbacks:** Handles `socket.on('close')` type events

#### Example

```
setImmediate(() => console.log('immediate'));
setTimeout(() => console.log('timeout'), 0);
Promise.resolve().then(() => console.log('promise'));
console.log('sync');
// Output: sync, promise, timeout, immediate
```

**Key Difference:** Node.js uses libuv for thread pool operations (file system, DNS), while browsers integrate with rendering cycles and have a simpler microtask/macrotask queue structure.

### 2. How would you implement optimistic UI updates in a React application with proper rollback on failure?

#### Optimistic UI Pattern

Optimistic UI immediately reflects user actions before server confirmation, improving perceived performance. Proper implementation requires **state management**, **rollback logic**, and **error handling**.

#### Implementation Strategy

```
const useOptimisticUpdate = (mutationFn) => {
  const [data, setData] = useState(initialData);

  const optimisticUpdate = async (newData) => {
    const previous = data;
    setData(newData);

    try {
      await mutationFn(newData);
    } catch (error) {
      setData(previous);
      showError(error);
    }
  };

  return [data, optimisticUpdate];
};
```

## Best Practices

- **Store previous state:** Always maintain a snapshot before optimistic update
- **Use unique IDs:** Generate temporary IDs for new items (e.g., temp-`{Date.now()}`)
- **Visual feedback:** Show pending states with opacity or loading indicators
- **Conflict resolution:** Handle cases where server returns different data
- **React Query/SWR:** Leverage libraries with built-in optimistic update support

### 3. What are the differences between SQL and NoSQL transaction guarantees, and when would you choose each?

#### SQL Transactions (ACID)

- **Atomicity:** All operations succeed or all fail
- **Consistency:** Database moves from one valid state to another
- **Isolation:** Concurrent transactions don't interfere
- **Durability:** Committed data persists through failures

#### NoSQL Transactions (BASE)

- **Basically Available:** System appears to work most of the time
- **Soft state:** State may change without input (eventual consistency)
- **Eventual consistency:** System becomes consistent over time

## Decision Matrix

### Choose SQL when:

- Complex relationships and joins are required
- Strong consistency is critical (financial transactions)
- ACID guarantees are mandatory
- Data structure is well-defined and stable

### Choose NoSQL when:

- Horizontal scaling is priority
- Schema flexibility is needed
- High write throughput is required
- Eventual consistency is acceptable (social media feeds, analytics)

**Modern Approach:** Many NoSQL databases (MongoDB, DynamoDB) now offer ACID transactions for specific use cases, blurring the distinction.

### 4. Explain how you would implement server-side rendering (SSR) with hydration in a React application and what challenges to watch for.

#### SSR with Hydration Process

**Server-side rendering** generates HTML on the server, while **hydration** attaches React event handlers to existing markup on the client.

#### Implementation Flow

```
// Server (Express + React)
app.get('*', async (req, res) => {
  const data = await fetchData(req.url);
  const html = renderToString(

  );
  res.send(`

  ${html}

  `);
});

// Client hydration
```

```
const data = window.__DATA__;
hydrateRoot(document.getElementById('root'), );
```

## Key Challenges

- **Hydration mismatch:** Server/client HTML must match exactly or React throws warnings
- **Data fetching:** Serialize initial data to avoid refetching on client
- **Browser-only APIs:** Check typeof window before using localStorage, document, etc.
- **Performance:** Large initial bundles delay interactivity (Time to Interactive)
- **State synchronization:** Ensure Redux/Context state matches between server and client

## Modern Solutions

Next.js and Remix handle these complexities with **automatic code splitting**, **streaming SSR**, and **progressive hydration**.

### 5. How do you handle database connection pooling in a production Node.js application, and what are the critical configuration parameters?

#### Connection Pooling Concept

Connection pooling maintains a **cache of database connections** that can be reused, avoiding the overhead of creating new connections for each query.

#### Implementation Example (PostgreSQL)

```
const { Pool } = require('pg');

const pool = new Pool({
  host: process.env.DB_HOST,
  database: process.env.DB_NAME,
  max: 20,
  idleTimeoutMillis: 30000,
  connectionTimeoutMillis: 2000,
});

const result = await pool.query('SELECT * FROM users WHERE id = $1', [userId]);
```

#### Critical Parameters

- **max (pool size):** Maximum connections (typically 10-20 per instance). Formula:  $(\text{core\_count} * 2) + \text{effective\_spindle\_count}$
- **min (minimum idle):** Connections kept alive (usually 2-5)
- **idleTimeoutMillis:** Time before closing idle connections (30000ms typical)
- **connectionTimeoutMillis:** Max wait time for available connection (2000-5000ms)
- **acquireTimeoutMillis:** Timeout for acquiring connection from pool

#### Production Best Practices

- Monitor active/idle connection ratio
- Implement proper error handling and connection release
- Use connection pool per database, not per request
- Configure based on database max\_connections limit
- Implement graceful shutdown to drain pool

### 6. Describe how JWT authentication works and what security considerations you must implement in a production system.

#### JWT Structure and Flow

JWT (JSON Web Token) consists of three parts: **Header** (algorithm), **Payload** (claims), and **Signature** (verification).

#### Authentication Flow

```
// Token generation
const token = jwt.sign(
```

```
{ userId: user.id, role: user.role },
process.env.JWT_SECRET,
{ expiresIn: '15m', algorithm: 'HS256' }
);
```

```
// Token verification
const decoded = jwt.verify(token, process.env.JWT_SECRET);
```

## Critical Security Considerations

- **Short expiration:** Access tokens should expire in 15-30 minutes
- **Refresh token pattern:** Use long-lived refresh tokens stored securely
- **Strong secrets:** Use cryptographically random secrets (256+ bits)
- **Algorithm specification:** Always specify algorithm to prevent 'none' attack
- **HTTPS only:** Never transmit tokens over unencrypted connections
- **HttpOnly cookies:** Store tokens in HttpOnly cookies to prevent XSS
- **CSRF protection:** Implement CSRF tokens when using cookies
- **Token revocation:** Maintain blacklist or use short-lived tokens

## Storage Strategy

**Access Token:** Memory or HttpOnly cookie

**Refresh Token:** HttpOnly, Secure, SameSite cookie with rotation

**7. What strategies would you use to optimize a React application that's experiencing performance issues with large lists?**

## Performance Optimization Strategies

### 1. Virtualization

Render only visible items using **react-window** or **react-virtualized**.

```
import { FixedSizeList } from 'react-window';
```

```
const Row = ({ index, style }) => (
```

```
Item {index}
```

```
);
```

```
{Row}
```

### 2. Memoization

- **React.memo:** Prevent re-renders of unchanged components
- **useMemo:** Cache expensive computations
- **useCallback:** Stabilize function references

### 3. Key Optimization

Use stable, unique keys (not array indices) to help React identify items efficiently.

### 4. Pagination/Infinite Scroll

Load data in chunks rather than all at once.

### 5. Web Workers

Offload heavy filtering/sorting to background threads.

### 6. Lazy Loading

```
const HeavyComponent = lazy(() => import('./Heavy'));
```

## Measurement Tools

- React DevTools Profiler
- Chrome Performance tab
- web-vitals library for Core Web Vitals

**8. Explain the differences between horizontal and vertical scaling, and how you would design a system to support horizontal scaling.**

## Scaling Approaches

**Vertical Scaling (Scale Up):** Adding more resources (CPU, RAM) to existing server

**Horizontal Scaling (Scale Out):** Adding more servers to distribute load

## Comparison

- **Vertical:** Simple, no code changes, but has hardware limits and single point of failure
- **Horizontal:** Unlimited scaling potential, fault tolerant, but requires architectural changes

## Designing for Horizontal Scaling

### 1. Stateless Application Servers

- Store session data in Redis/external cache, not server memory
- Any server should handle any request

### 2. Load Balancing

```
// Nginx load balancer config
upstream backend {
    server app1.example.com;
    server app2.example.com;
    server app3.example.com;
}
```

```
server {
    location / {
        proxy_pass http://backend;
    }
}
```

### 3. Database Considerations

- Read replicas for read-heavy workloads
- Sharding for write scaling
- Caching layer (Redis) to reduce database load

### 4. Shared Storage

- Use S3/object storage for files, not local filesystem
- Centralized logging and monitoring

### 5. Service Discovery

- Dynamic service registration (Consul, etcd)
- Health checks and automatic failover

**9. How would you implement rate limiting in a REST API, and what algorithms would you consider?**

## Rate Limiting Algorithms

### 1. Token Bucket

Most flexible. Tokens added at fixed rate, consumed per request. Allows bursts.

```
class TokenBucket {
```

```

constructor(capacity, refillRate) {
  this.capacity = capacity;
  this.tokens = capacity;
  this.refillRate = refillRate;
  this.lastRefill = Date.now();
}

consume() {
  this.refill();
  if (this.tokens > 0) {
    this.tokens--;
    return true;
  }
  return false;
}
}

```

## 2. Leaky Bucket

Requests processed at constant rate. Smooths traffic but less flexible.

## 3. Fixed Window

Simple counter reset at fixed intervals. Can allow 2x limit at window boundaries.

## 4. Sliding Window Log

Tracks timestamp of each request. Accurate but memory-intensive.

## Production Implementation

```

// Express middleware with Redis
const rateLimit = require('express-rate-limit');
const RedisStore = require('rate-limit-redis');

const limiter = rateLimit({
  store: new RedisStore({ client: redisClient }),
  windowMs: 15 * 60 * 1000,
  max: 100,
  message: 'Too many requests'
});

```

## Best Practices

- Return 429 status with Retry-After header
- Different limits per endpoint/user tier
- Implement at API gateway level
- Monitor and alert on rate limit hits
- Use distributed rate limiting (Redis) for multiple servers

## 10. What are the key differences between REST and GraphQL APIs, and in what scenarios would you choose one over the other?

### Architecture Comparison

**REST:** Multiple endpoints, fixed data structures

**GraphQL:** Single endpoint, flexible queries

### Key Differences

- **Data Fetching:** REST may over-fetch/under-fetch; GraphQL requests exact data needed
- **Versioning:** REST uses v1, v2 endpoints; GraphQL evolves schema with deprecation
- **Caching:** REST uses HTTP caching; GraphQL requires custom caching strategies
- **Learning Curve:** REST is simpler; GraphQL has steeper initial learning

## GraphQL Example

```
// Query exact fields needed
query {
  user(id: "123") {
    name
    posts(limit: 5) {
      title
      comments { author }
    }
  }
}
```

## **When to Choose REST**

- Simple, resource-based operations (CRUD)
- Heavy caching requirements
- File uploads/downloads
- Team unfamiliar with GraphQL
- Public APIs requiring broad compatibility

## **When to Choose GraphQL**

- Complex, nested data relationships
- Multiple client types (web, mobile) with different needs
- Rapidly evolving frontend requirements
- Need to aggregate data from multiple sources
- Want strong typing and introspection

## **Hybrid Approach**

Many organizations use REST for simple operations and GraphQL for complex queries, or implement GraphQL as a gateway over REST microservices.

## Data Structures and Algorithms

Questions in this section test your understanding of how to work with and manipulate data efficiently.

---

### 1. Explain how to implement an LRU (Least Recently Used) Cache with $O(1)$ time complexity for both get and put operations.

#### LRU Cache Implementation

An **LRU Cache** requires a combination of a **HashMap** and a **Doubly Linked List**. The HashMap provides  $O(1)$  lookup, while the doubly linked list maintains the order of usage.

- **HashMap:** Maps keys to nodes in the linked list
- **Doubly Linked List:** Maintains access order (most recent at head, least recent at tail)
- **Get operation:** Move accessed node to head
- **Put operation:** Add new node at head, remove tail if capacity exceeded

```
class LRUCache {
  constructor(capacity) {
    this.capacity = capacity;
    this.cache = new Map();
    this.head = { prev: null, next: null };
    this.tail = { prev: this.head, next: null };
    this.head.next = this.tail;
  }
}
```

**Time Complexity:**  $O(1)$  for both get and put operations

**Space Complexity:**  $O(\text{capacity})$

### 2. How would you find all pairs in an array that sum to a target value? What's the optimal approach?

#### Two Sum / Pair Sum Problem

The optimal approach uses a **HashSet** to achieve  $O(n)$  time complexity with a single pass through the array.

#### Algorithm:

- Iterate through the array once
- For each element, check if  $(\text{target} - \text{element})$  exists in the set
- If found, record the pair
- Add current element to the set

```
function findPairs(arr, target) {
  const seen = new Set();
  const pairs = [];
  for (let num of arr) {
    if (seen.has(target - num)) {
      pairs.push([target - num, num]);
    }
    seen.add(num);
  }
  return pairs;
}
```

**Time Complexity:**  $O(n)$

**Space Complexity:**  $O(n)$

**3. Implement a function to detect a cycle in a linked list. Explain the algorithm and its complexity.**

### Floyd's Cycle Detection Algorithm

Use the **two-pointer technique** (tortoise and hare) where one pointer moves twice as fast as the other.

#### How it works:

- **Slow pointer:** Moves one step at a time
- **Fast pointer:** Moves two steps at a time
- If there's a cycle, fast will eventually meet slow
- If fast reaches null, there's no cycle

```
function hasCycle(head) {
  let slow = head, fast = head;
  while (fast && fast.next) {
    slow = slow.next;
    fast = fast.next.next;
    if (slow === fast) return true;
  }
  return false;
}
```

**Time Complexity:**  $O(n)$

**Space Complexity:**  $O(1)$

**4. What is the sliding window technique? Provide an example of finding the maximum sum of a subarray of size k.**

### Sliding Window Technique

The **sliding window** is an optimization technique for problems involving contiguous subarrays or substrings. Instead of recalculating from scratch, we slide the window by removing one element and adding another.

#### Maximum Sum Subarray of Size K:

- Calculate sum of first k elements
- Slide window: subtract leftmost, add new rightmost
- Track maximum sum encountered

```
function maxSumSubarray(arr, k) {
  let maxSum = 0, windowSum = 0;
  for (let i = 0; i < k; i++) windowSum += arr[i];
  maxSum = windowSum;
  for (let i = k; i < arr.length; i++) {
    windowSum = windowSum - arr[i - k] + arr[i];
    maxSum = Math.max(maxSum, windowSum);
  }
  return maxSum;
}
```

**Time Complexity:**  $O(n)$  vs  $O(n*k)$  brute force

**5. Explain the difference between a Stack and a Queue. How would you implement a Queue using two Stacks?**

### Stack vs Queue

- **Stack:** LIFO (Last In First Out) - push/pop from same end
- **Queue:** FIFO (First In First Out) - enqueue at rear, dequeue from front

#### Queue Using Two Stacks:

Use two stacks: **inbox** (for enqueue) and **outbox** (for dequeue). When dequeuing, if outbox is

empty, transfer all elements from inbox to outbox, reversing their order.

```
class QueueWithStacks {
  constructor() {
    this.inbox = [];
    this.outbox = [];
  }
  enqueue(x) { this.inbox.push(x); }
  dequeue() {
    if (!this.outbox.length)
      while (this.inbox.length) this.outbox.push(this.inbox.pop());
    return this.outbox.pop();
  }
}
```

**Time Complexity:** Amortized  $O(1)$  for both operations

## 6. What is a Trie data structure and when would you use it? Implement basic insert and search operations.

### Trie (Prefix Tree)

A **Trie** is a tree-like data structure for storing strings where each node represents a character. It's highly efficient for prefix-based operations.

#### Use Cases:

- Autocomplete systems
- Spell checkers
- IP routing tables
- Dictionary implementations

```
class TrieNode {
  constructor() {
    this.children = {};
    this.isEndOfWord = false;
  }
}
class Trie {
  constructor() { this.root = new TrieNode(); }
  insert(word) {
    let node = this.root;
    for (let char of word) {
      if (!node.children[char]) node.children[char] = new TrieNode();
      node = node.children[char];
    }
    node.isEndOfWord = true;
  }
  search(word) {
    let node = this.root;
    for (let char of word) {
      if (!node.children[char]) return false;
      node = node.children[char];
    }
    return node.isEndOfWord;
  }
}
```

**Time Complexity:**  $O(m)$  where  $m$  is word length

## 7. Explain how a HashMap works internally. What happens during collisions and how does resizing work?

### HashMap Internal Implementation

A **HashMap** uses an array of buckets where each bucket can store key-value pairs. The index is determined by hashing the key.

## Key Concepts:

- **Hash Function:** Converts key to array index:  $\text{index} = \text{hash}(\text{key}) \% \text{arraySize}$
- **Collisions:** When two keys hash to same index
- **Collision Resolution:** Chaining (linked list) or Open Addressing (probing)
- **Load Factor:** ratio of entries to buckets (typically 0.75)
- **Resizing:** When load factor exceeded, double array size and rehash all entries

## Chaining Example:

```
class HashMap {
  constructor() {
    this.buckets = new Array(16).fill(null).map(() => []);
    this.size = 0;
  }
  hash(key) {
    return key.toString().split('').reduce((acc, char) =>
      acc + char.charCodeAt(0), 0) % this.buckets.length;
  }
}
```

**Time Complexity:** Average  $O(1)$ , Worst  $O(n)$

## 8. What is the difference between BFS and DFS? When would you choose one over the other?

### BFS vs DFS

**Breadth-First Search (BFS):** Explores level by level using a queue

**Depth-First Search (DFS):** Explores as deep as possible using a stack or recursion

### When to Use:

- **BFS:** Shortest path in unweighted graphs, level-order traversal, finding nearest neighbor
- **DFS:** Topological sorting, cycle detection, path finding, maze solving, tree traversals

## BFS Implementation:

```
function bfs(graph, start) {
  const visited = new Set();
  const queue = [start];
  visited.add(start);
  while (queue.length) {
    const node = queue.shift();
    for (let neighbor of graph[node]) {
      if (!visited.has(neighbor)) {
        visited.add(neighbor);
        queue.push(neighbor);
      }
    }
  }
}
```

**Space Complexity:** BFS  $O(w)$  width, DFS  $O(h)$  height

## 9. Implement a function to merge K sorted linked lists efficiently. What's the optimal time complexity?

### Merge K Sorted Lists

The optimal approach uses a **Min Heap (Priority Queue)** to efficiently track the smallest element among all list heads.

### Algorithm:

- Insert the head of each list into a min heap
- Extract minimum, add to result

- Insert the next node from that list into heap
- Repeat until heap is empty

```
function mergeKLists(lists) {
  const heap = new MinHeap();
  for (let list of lists) {
    if (list) heap.insert(list);
  }
  const dummy = { next: null };
  let current = dummy;
  while (!heap.isEmpty()) {
    const node = heap.extractMin();
    current.next = node;
    current = current.next;
    if (node.next) heap.insert(node.next);
  }
  return dummy.next;
}
```

**Time Complexity:**  $O(N \log k)$  where  $N$  is total nodes,  $k$  is number of lists

**Space Complexity:**  $O(k)$

## 10. What is Dynamic Programming? Explain with the example of calculating Fibonacci numbers.

### Dynamic Programming (DP)

**Dynamic Programming** is an optimization technique that solves complex problems by breaking them into overlapping subproblems and storing results to avoid redundant calculations.

#### Two Approaches:

- **Memoization (Top-Down):** Recursive with caching
- **Tabulation (Bottom-Up):** Iterative with array

#### Fibonacci - Memoization:

```
function fib(n, memo = {}) {
  if (n <= 1) return n;
  if (memo[n]) return memo[n];
  memo[n] = fib(n - 1, memo) + fib(n - 2, memo);
  return memo[n];
}
```

#### Fibonacci - Tabulation:

```
function fib(n) {
  if (n <= 1) return n;
  const dp = [0, 1];
  for (let i = 2; i <= n; i++) {
    dp[i] = dp[i - 1] + dp[i - 2];
  }
  return dp[n];
}
```

**Time:**  $O(n)$ , **Space:**  $O(n)$  or  $O(1)$  with optimization

## System Design

These questions evaluate your ability to think about the bigger picture, including architecture, scalability, and performance.

---

### 1. Design a scalable URL shortener service like bit.ly. What are the key components and how would you handle high traffic?

#### Key Components

- **URL Generation Service:** Creates short codes using base62 encoding or hash functions
- **Database:** Stores mapping between short and long URLs
- **Cache Layer:** Redis/Memcached for frequently accessed URLs
- **Load Balancer:** Distributes traffic across multiple servers
- **Analytics Service:** Tracks clicks and metrics

#### Architecture Considerations

- **Stateless servers:** Any server can handle any request
- **Database sharding:** Partition by hash of short URL for horizontal scaling
- **Read-heavy optimization:** 100:1 read-to-write ratio, use read replicas
- **Caching strategy:** Cache hot URLs with LRU eviction
- **Unique ID generation:** Use distributed ID generator (Snowflake) or pre-generated key pools

#### Sample Code for Base62 Encoding

```
function encodeBase62(num) {
  const chars = '0123456789abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ';
  let result = '';
  while (num > 0) {
    result = chars[num % 62] + result;
    num = Math.floor(num / 62);
  }
  return result || '0';
}
```

#### Scale Estimates

**Traffic:** 500M URLs/month = ~200 writes/sec, 20K reads/sec. **Storage:** 500 bytes per URL × 500M = 250GB/month.

### 2. How would you design a real-time chat application that supports millions of concurrent users?

#### Architecture Components

- **WebSocket servers:** Maintain persistent connections with clients
- **Message broker:** Kafka or RabbitMQ for message queuing
- **Presence service:** Tracks online/offline status using Redis
- **Message storage:** Cassandra or MongoDB for chat history
- **API Gateway:** REST APIs for authentication and user management
- **CDN:** Serve static assets and media files

#### Key Design Decisions

- **Connection management:** Each WebSocket server handles 50K-100K connections
- **Message delivery:** At-least-once delivery with idempotency keys
- **Sharding strategy:** Partition users by user\_id or chat\_room\_id
- **Horizontal scaling:** Stateless WebSocket servers behind load balancer
- **Message ordering:** Use sequence numbers and vector clocks

## WebSocket Connection Example

```
const WebSocket = require('ws');
const server = new WebSocket.Server({ port: 8080 });

server.on('connection', (socket) => {
  socket.on('message', (data) => {
    const msg = JSON.parse(data);
    broadcastToRoom(msg.roomId, msg);
  });
});
```

## Challenges

**Presence sync:** Use heartbeats every 30s. **Offline messages:** Queue in message broker until user reconnects.

**3. Design a social media news feed system like Twitter or Facebook. How do you ensure fast feed generation for users with millions of followers?**

## Two Main Approaches

- **Pull Model (Fan-out on read):** Generate feed when user requests it by querying followed users' posts
- **Push Model (Fan-out on write):** Pre-compute feed and push to followers' timelines when post is created
- **Hybrid Model:** Push for most users, pull for celebrities with millions of followers

## Recommended Hybrid Architecture

- **Post service:** Handles post creation and storage
- **Fan-out service:** Pushes posts to followers' feed caches
- **Feed service:** Aggregates and ranks posts for display
- **Graph database:** Stores follower relationships
- **Redis cache:** Stores pre-computed feeds (last 1000 posts)
- **Ranking algorithm:** ML-based relevance scoring

## Feed Generation Logic

```
async function getFeed(userId, limit) {
  const cached = await redis.get(`feed:${userId}`);
  if (cached) return JSON.parse(cached).slice(0, limit);

  const following = await getFollowing(userId);
  const posts = await fetchRecentPosts(following);
  const ranked = rankByRelevance(posts, userId);
  await redis.setex(`feed:${userId}`, 3600, JSON.stringify(ranked));
  return ranked.slice(0, limit);
}
```

## Scale Considerations

**Write amplification:** For users with 10M followers, queue fan-out jobs. **Read optimization:** Cache feeds with 1-hour TTL.

**4. Explain the CAP theorem and how it applies to distributed database design. Give examples of systems that prioritize different properties.**

## CAP Theorem Overview

**CAP theorem** states that a distributed system can only guarantee two of three properties:

- **Consistency (C):** All nodes see the same data at the same time
- **Availability (A):** Every request receives a response (success or failure)
- **Partition Tolerance (P):** System continues operating despite network partitions

## Practical Reality

Since network partitions are inevitable in distributed systems, you must choose between **CP (Consistency + Partition Tolerance)** or **AP (Availability + Partition Tolerance)**.

## CP Systems (Consistency over Availability)

- **HBase, MongoDB (with majority writes):** Block writes during partition to maintain consistency
- **Use case:** Financial transactions, inventory management
- **Behavior:** Returns error if consistency cannot be guaranteed

## AP Systems (Availability over Consistency)

- **Cassandra, DynamoDB, Riak:** Accept writes on any node, eventual consistency
- **Use case:** Social media feeds, shopping carts, analytics
- **Behavior:** Always accepts reads/writes, resolves conflicts later

## Tunable Consistency Example (Cassandra)

```
// Strong consistency: R + W > N
SELECT * FROM users WHERE id = 123
  USING CONSISTENCY QUORUM;
```

```
// Eventual consistency
SELECT * FROM users WHERE id = 123
  USING CONSISTENCY ONE;
```

## Modern Perspective

Systems like **Spanner** achieve strong consistency globally using atomic clocks and GPS.

**5. Design a distributed rate limiter that can handle 100K requests per second. What algorithms and data structures would you use?**

## Rate Limiting Algorithms

- **Token Bucket:** Tokens added at fixed rate, requests consume tokens
- **Leaky Bucket:** Requests processed at fixed rate, excess queued or dropped
- **Fixed Window:** Count requests in fixed time windows (prone to burst at boundaries)
- **Sliding Window Log:** Track timestamp of each request
- **Sliding Window Counter:** Weighted count from current and previous window

## Recommended Architecture

- **Redis cluster:** Centralized rate limit counters with atomic operations
- **Token bucket algorithm:** Best balance of smoothness and burst handling
- **Distributed cache:** Partition by user\_id or API\_key for horizontal scaling
- **Local cache:** In-memory cache with sync to Redis every second

## Redis Token Bucket Implementation

```
function allowRequest(userId, maxTokens, refillRate) {
  const key = `rate_limit:${userId}`;
  const now = Date.now();
  const [tokens, lastRefill] = redis.hmget(key, 'tokens', 'last');
  const elapsed = (now - lastRefill) / 1000;
  const newTokens = Math.min(maxTokens, tokens + elapsed * refillRate);
  if (newTokens >= 1) {
    redis.hmset(key, 'tokens', newTokens - 1, 'last', now);
    return true;
  }
  return false;
}
```

## Scale Considerations

**Partitioning:** Shard Redis by consistent hashing. **Performance:** Use Lua scripts for atomic operations. **Fallback:** Allow requests if Redis is down (fail-open).

## 6. How would you design a notification system that supports multiple channels (email, SMS, push) and handles millions of notifications per day?

### System Components

- **API Gateway:** Receives notification requests from services
- **Message Queue:** Kafka or SQS for buffering and decoupling
- **Notification Service:** Routes to appropriate channel handlers
- **Channel Workers:** Dedicated workers for email, SMS, push
- **Template Service:** Manages notification templates and personalization
- **User Preference Service:** Stores opt-in/opt-out preferences
- **Analytics Service:** Tracks delivery, opens, clicks

### Architecture Patterns

- **Priority queues:** Separate queues for urgent vs. bulk notifications
- **Rate limiting:** Respect provider limits (Twilio, SendGrid)
- **Retry logic:** Exponential backoff for failed deliveries
- **Idempotency:** Deduplicate notifications using unique IDs
- **Circuit breaker:** Prevent cascading failures

### Notification Worker Example

```
async function processNotification(msg) {
  const { userId, type, channel, template } = msg;
  const prefs = await getUserPreferences(userId);
  if (!prefs[channel]) return;

  const content = await renderTemplate(template, msg.data);
  const result = await sendViaChannel(channel, userId, content);
  await logDelivery(msg.id, result);
}
```

### Scaling Strategy

**Throughput:** 1M notifications/day = ~12 per second. Use worker pools with auto-scaling. **Cost optimization:** Batch emails, use cheaper channels when possible.

## 7. Design a file storage system like Dropbox or Google Drive. How do you handle file synchronization, versioning, and conflict resolution?

### Core Components

- **Metadata service:** Stores file metadata, hierarchy, permissions
- **Block storage:** S3 or distributed file system for actual file content
- **Sync service:** Detects changes and propagates updates
- **Notification service:** WebSocket for real-time sync notifications
- **Chunking service:** Splits files into blocks for efficient transfer
- **Version control:** Maintains file history and snapshots

### Key Design Decisions

- **Chunking strategy:** Fixed-size (4MB) or content-defined (CDC) chunks
- **Deduplication:** Hash-based deduplication at chunk level
- **Delta sync:** Transfer only changed chunks, not entire files
- **Conflict resolution:** Last-write-wins or manual merge for simultaneous edits
- **Offline support:** Queue operations locally, sync when online

### File Sync Algorithm

```
async function syncFile(fileId, localHash, localVersion) {
  const remote = await getFileMetadata(fileId);
  if (remote.version > localVersion) {
    const chunks = await getModifiedChunks(fileId, localVersion);
    await downloadAndMerge(chunks);
  } else if (localHash !== remote.hash) {
    const diff = await computeDiff(fileId);
  }
}
```

```
    await uploadChunks(diff);
  }
}
```

## Versioning Strategy

**Metadata:** Store version tree with parent pointers. **Storage optimization:** Keep only deltas between versions. **Retention:** 30-day version history for free users.

**8. Explain the difference between stateless and stateful architectures. When would you choose one over the other?**

### Stateless Architecture

- **Definition:** Server does not store client session data between requests
- **State storage:** Client-side (JWT tokens) or external store (Redis)
- **Advantages:** Easy horizontal scaling, fault tolerance, simple load balancing
- **Disadvantages:** Larger request payloads, potential security concerns
- **Use cases:** RESTful APIs, microservices, serverless functions

### Stateful Architecture

- **Definition:** Server maintains session state in memory or local storage
- **State storage:** In-memory on specific server instance
- **Advantages:** Lower latency, smaller payloads, easier session management
- **Disadvantages:** Sticky sessions required, harder to scale, single point of failure
- **Use cases:** WebSocket connections, gaming servers, real-time collaboration

### Stateless API Example (JWT)

```
app.get('/api/profile', authenticate, (req, res) => {
  // Token decoded in middleware, no server-side session
  const userId = req.user.id;
  const profile = await db.getProfile(userId);
  res.json(profile);
});
```

### Hybrid Approach

- **Stateless for APIs:** Business logic services remain stateless
- **Stateful for connections:** WebSocket/SSE servers maintain connections
- **External state store:** Redis for sessions shared across servers

### Decision Criteria

**Choose stateless:** High scalability needs, cloud-native apps. **Choose stateful:** Real-time features, performance-critical applications.

**9. Design a distributed caching system. How would you handle cache invalidation, consistency, and eviction policies?**

### Architecture Components

- **Cache layer:** Redis cluster or Memcached for distributed caching
- **Cache client:** Application-level client with consistent hashing
- **Write-through/Write-back:** Caching strategies for writes
- **CDN layer:** Edge caching for static content
- **Local cache:** In-memory L1 cache for hot data

### Cache Invalidation Strategies

- **Time-based (TTL):** Expire cache after fixed duration
- **Event-based:** Invalidate on data updates using pub/sub
- **Write-through:** Update cache synchronously with database
- **Cache-aside:** Application manages cache population
- **Write-behind:** Async write to DB, immediate cache update

## Cache-Aside Pattern Implementation

```
async function getData(key) {
  let data = await cache.get(key);
  if (!data) {
    data = await db.query(key);
    await cache.setex(key, 3600, JSON.stringify(data));
  }
  return data;
}
```

```
async function updateData(key, value) {
  await db.update(key, value);
  await cache.del(key);
}
```

## Eviction Policies

- **LRU (Least Recently Used):** Most common, evicts least accessed items
- **LFU (Least Frequently Used):** Tracks access frequency
- **FIFO:** First in, first out
- **TTL-based:** Automatic expiration

## Consistency Challenges

**Cache stampede:** Use locking or probabilistic early expiration. **Thundering herd:** Stagger TTLs. **Stale data:** Accept eventual consistency or use versioning.

**10. How would you design a search engine like Elasticsearch for an e-commerce platform? Include indexing, ranking, and query optimization.**

## System Architecture

- **Indexing service:** Processes product catalog and builds inverted index
- **Search cluster:** Distributed Elasticsearch/Solr nodes
- **Query parser:** Analyzes and optimizes search queries
- **Ranking service:** Scores results based on relevance and business rules
- **Auto-complete service:** Suggests queries using trie data structure
- **Analytics service:** Tracks search metrics and user behavior

## Indexing Strategy

- **Document structure:** Flatten product attributes for efficient querying
- **Analyzers:** Tokenization, stemming, stop words, synonyms
- **Sharding:** Partition index by category or hash for scalability
- **Replication:** Multiple replicas for high availability
- **Update strategy:** Near real-time indexing with refresh intervals

## Elasticsearch Query Example

```
{
  "query": {
    "bool": {
      "must": [{"match": {"title": "laptop"}}],
      "filter": [
        {"range": {"price": {"lte": 1000}}},
        {"term": {"brand": "dell"}}
      ]
    }
  },
  "sort": [{"_score": "desc"}, {"price": "asc"}]
}
```

## Ranking Factors

- **Text relevance:** TF-IDF or BM25 scoring
- **Popularity:** Sales volume, ratings, reviews

- **Personalization:** User history and preferences
- **Business rules:** Promote featured products, margins

## **Performance Optimization**

**Caching:** Cache popular queries. **Filters:** Use filters over queries when possible. **Pagination:** Use `search_after` for deep pagination.

## Coding and Debugging

This section presents practical coding challenges and questions about debugging techniques.

### 1. Write a function to flatten a nested array of arbitrary depth without using built-in flatten methods.

#### Solution

Here's an efficient recursive approach to flatten nested arrays:

```
function flattenArray(arr) {
  const result = [];
  for (let item of arr) {
    if (Array.isArray(item)) {
      result.push(...flattenArray(item));
    } else {
      result.push(item);
    }
  }
  return result;
}
```

#### Key points:

- Uses recursion to handle arbitrary nesting depth
- Checks each element with `Array.isArray()`
- Spread operator merges nested results efficiently
- Time complexity:  $O(n)$  where  $n$  is total number of elements

### 2. Implement a function to check if a string is a palindrome, considering only alphanumeric characters and ignoring case.

#### Solution

Here's an optimized two-pointer approach:

```
function isPalindrome(str) {
  const cleaned = str.toLowerCase().replace(/[^a-z0-9]/g, "");
  let left = 0, right = cleaned.length - 1;
  while (left < right) {
    if (cleaned[left] !== cleaned[right]) return false;
    left++; right--;
  }
  return true;
}
```

#### Explanation:

- First normalize the string by converting to lowercase and removing non-alphanumeric characters
- Use two pointers from both ends moving toward center
- Time complexity:  $O(n)$ , Space complexity:  $O(n)$  for cleaned string
- Alternative: Check characters in-place without creating new string for  $O(1)$  space

### 3. What are the key differences between debugging techniques in development vs production environments?

#### Development Debugging

- **Console logging:** Liberal use of `console.log`, debugger statements
- **Source maps:** Unminified code with full stack traces

- **Hot reloading:** Immediate feedback on code changes
- **Browser DevTools:** Breakpoints, step-through debugging, React/Vue DevTools

## Production Debugging

- **Error monitoring:** Tools like Sentry, Rollbar, or New Relic for centralized error tracking
- **Structured logging:** Winston, Bunyan with log levels and correlation IDs
- **APM tools:** Application Performance Monitoring for bottleneck identification
- **Feature flags:** Enable/disable features without deployment
- **Remote debugging:** Limited access, relies on logs and metrics rather than live debugging

**Best practice:** Never leave console.log or debugger statements in production code.

## 4. Explain memory profiling techniques and how to identify memory leaks in a full stack application.

### Frontend Memory Profiling

- **Chrome DevTools Memory Profiler:** Heap snapshots to compare memory usage over time
- **Performance Monitor:** Track DOM nodes, event listeners, and JS heap size in real-time
- **Common causes:** Detached DOM nodes, global variables, event listeners not removed, closures holding references

### Backend Memory Profiling

- **Node.js:** Use --inspect flag with Chrome DevTools or clinic.js for diagnostics
- **Heap dumps:** Generate with process.memoryUsage() or heapdump module
- **Common causes:** Unclosed database connections, growing caches, circular references, stream leaks

### Detection Strategy

- Monitor memory usage trends over time
- Take heap snapshots before and after operations
- Look for objects that aren't garbage collected
- Use weak references (WeakMap, WeakSet) where appropriate

## 5. Write a debounce function from scratch and explain when you'd use it in a full stack context.

### Implementation

```
function debounce(func, delay) {
  let timeoutId;
  return function(...args) {
    clearTimeout(timeoutId);
    timeoutId = setTimeout(() => {
      func.apply(this, args);
    }, delay);
  };
}
```

### Use Cases

- **Search autocomplete:** Wait for user to stop typing before API call
- **Window resize handlers:** Delay expensive layout recalculations
- **Form validation:** Validate after user pauses input
- **Scroll events:** Reduce frequency of scroll-triggered operations
- **API rate limiting:** Prevent excessive backend requests

**Key difference from throttle:** Debounce delays execution until after the event stops firing, while throttle ensures execution at regular intervals.

## 6. How do you handle and debug race conditions in asynchronous code?

### Common Race Condition Scenarios

- **Multiple API calls:** Results arrive out of order
- **State updates:** Concurrent setState calls in React
- **Database transactions:** Simultaneous reads/writes

## Prevention Techniques

```
// Use request cancellation
let currentRequest = null;
async function fetchData(query) {
  currentRequest?.abort();
  const controller = new AbortController();
  currentRequest = controller;
  const response = await fetch(url, {signal: controller.signal});
  return response.json();
}
```

- **Request tokens:** Track latest request with incrementing ID
- **Promises with flags:** Check if component is still mounted
- **Database locks:** Use optimistic/pessimistic locking strategies
- **Atomic operations:** Use transactions with isolation levels
- **Message queues:** Serialize operations through queue systems

## 7. What is monkey patching and when is it appropriate to use? Provide an example.

### Definition

**Monkey patching** is dynamically modifying or extending code at runtime, typically by adding/replacing methods on classes or objects.

### Example

```
// Patching Array prototype
Array.prototype.last = function() {
  return this[this.length - 1];
};

const arr = [1, 2, 3];
console.log(arr.last()); // 3
```

### Appropriate Use Cases

- **Polyfills:** Adding missing browser features for compatibility
- **Testing:** Mocking dependencies or system functions
- **Hotfixes:** Temporary fixes for third-party library bugs
- **Instrumentation:** Adding logging or monitoring to existing code

### Risks and Alternatives

- Can cause conflicts with other code or future language features
- Makes debugging harder and code less predictable
- **Better alternatives:** Wrapper functions, composition, decorator pattern, or proper inheritance

## 8. Explain exception handling best practices across the full stack. How do you structure try-catch blocks?

### Frontend Exception Handling

- **Global error boundary:** React Error Boundaries to catch component errors
- **Window error handler:** window.onerror and window.onunhandledrejection
- **Async/await:** Always wrap in try-catch for user-facing operations

```
async function fetchUser(id) {
  try {
    const response = await api.getUser(id);
    return response.data;
  } catch (error) {
    logger.error('Failed to fetch user', {id, error});
    throw new UserFetchError('Unable to load user data');
  }
}
```

```
}  
}
```

## Backend Best Practices

- **Centralized error middleware:** Express error handlers for consistent responses
- **Custom error classes:** Create specific error types (ValidationError, AuthError)
- **Fail fast:** Validate inputs early, throw meaningful errors
- **Log with context:** Include request IDs, user info, stack traces
- **Never expose internals:** Return sanitized errors to clients

### 9. Write a function to deep clone an object, handling circular references.

#### Solution

```
function deepClone(obj, hash = new WeakMap()) {  
  if (obj === null || typeof obj !== 'object') return obj;  
  if (hash.has(obj)) return hash.get(obj);  
  
  const clone = Array.isArray(obj) ? [] : {};  
  hash.set(obj, clone);  
  
  for (let key in obj) {  
    if (obj.hasOwnProperty(key)) clone[key] = deepClone(obj[key], hash);  
  }  
  return clone;  
}
```

#### Key Features

- **WeakMap tracking:** Prevents infinite loops from circular references
- **Type checking:** Handles primitives, arrays, and objects
- **Recursive cloning:** Deep copies nested structures
- Preserves array vs object distinction

**Limitations:** Doesn't handle functions, Dates, RegExp, Maps, Sets. For complete cloning, use structuredClone() API or libraries like lodash.cloneDeep.

### 10. What debugging tools and techniques do you use for diagnosing performance bottlenecks in a full stack application?

#### Frontend Performance Tools

- **Chrome DevTools Performance tab:** Record runtime performance, analyze frame rates, identify long tasks
- **Lighthouse:** Automated audits for performance, accessibility, SEO
- **React Profiler:** Identify unnecessary re-renders and component performance
- **Network tab:** Analyze request timing, payload sizes, waterfall charts
- **Coverage tool:** Find unused CSS and JavaScript

#### Backend Performance Tools

- **APM solutions:** New Relic, DataDog for end-to-end tracing
- **Database query profiling:** EXPLAIN ANALYZE in PostgreSQL, slow query logs
- **Node.js profiling:** 0x, clinic.js flame graphs for CPU profiling
- **Load testing:** Artillery, k6 for stress testing under load

#### Key Metrics

- Time to First Byte (TTFB), First Contentful Paint (FCP), Time to Interactive (TTI)
- Database query execution time and N+1 query detection
- API response times and throughput

## Behavioral Questions

These questions assess your soft skills, problem-solving approach, and how you work in a team.

---

### 1. Tell me about a time when you had to debug a critical production issue under pressure.

**Situation:** Our e-commerce platform experienced a sudden 500% increase in database query times during Black Friday, causing checkout failures and revenue loss of \$10K per minute.

**Task:** As the lead full stack developer, I needed to identify and resolve the issue within 30 minutes to minimize business impact.

**Action:** I immediately checked application logs and APM tools, identified an unoptimized N+1 query in the order processing service. I implemented query batching and added database indexes, then deployed a hotfix using our blue-green deployment strategy.

**Result:** Query times dropped from 8 seconds to 200ms within 15 minutes. We recovered the checkout flow and prevented an estimated \$500K in lost sales. I later conducted a post-mortem and implemented query performance monitoring alerts.

### 2. Describe a situation where you had to make a difficult technical decision that involved trade-offs.

**Situation:** Our SaaS application needed real-time notifications, and I had to choose between WebSockets, Server-Sent Events (SSE), or polling. Each had different implications for scalability, cost, and browser compatibility.

**Task:** Select the optimal solution that balanced real-time performance with infrastructure costs while supporting 100K+ concurrent users.

**Action:** I created a comparison matrix evaluating connection overhead, server resources, fallback mechanisms, and scaling costs. I built POC implementations for WebSockets and SSE, conducted load testing, and presented findings to stakeholders. We chose SSE for its simplicity, automatic reconnection, and 40% lower server costs compared to WebSockets.

**Result:** Successfully deployed SSE-based notifications handling 150K concurrent connections with 99.9% uptime. The solution saved \$50K annually in infrastructure costs while meeting all real-time requirements.

### 3. Give an example of when you had to mentor or help a junior developer overcome a technical challenge.

**Situation:** A junior developer on my team was struggling with implementing authentication using JWT tokens and kept creating security vulnerabilities like storing tokens in localStorage and not validating token expiration properly.

**Task:** Help them understand secure authentication practices while building their confidence and problem-solving skills.

**Action:** I scheduled pair programming sessions where we refactored their code together. I explained security concepts like XSS attacks, token storage best practices using httpOnly cookies, and implementing refresh token rotation. I shared resources on OWASP guidelines and created a code review checklist for authentication flows.

**Result:** The developer successfully implemented secure authentication that passed security audit. They became our team's go-to person for auth-related questions and later presented a tech talk on JWT security to the entire engineering team.

### 4. Tell me about a time when you disagreed with a technical approach proposed by your team or manager.

**Situation:** Our product manager wanted to implement a microservices architecture for a new

feature, but our application was a well-structured monolith serving only 5,000 users. I believed microservices would introduce unnecessary complexity.

**Task:** Advocate for a pragmatic solution while respecting the PM's concerns about future scalability.

**Action:** I prepared a technical proposal comparing both approaches with metrics on development time, operational overhead, and scaling thresholds. I suggested a modular monolith with clear service boundaries that could be extracted later if needed. I presented data showing we wouldn't need microservices until reaching 50K+ users based on our current architecture's capacity.

**Result:** The team agreed to start with the modular monolith approach, reducing initial development time by 6 weeks. When we reached 40K users 18 months later, we successfully extracted two services with minimal refactoring because of the clear boundaries we had established.

## **5. Describe a situation where you had to learn a new technology quickly to complete a project.**

**Situation:** Our client required integrating real-time video processing capabilities into our application within 3 weeks, but none of our team had experience with WebRTC or media streaming protocols.

**Task:** Learn WebRTC, implement peer-to-peer video connections, and build a scalable signaling server as the technical lead.

**Action:** I dedicated 2 days to intensive learning through documentation, video tutorials, and building small prototypes. I experimented with libraries like SimpleWebRTC and PeerJS, evaluated STUN/TURN server options, and created a proof-of-concept. I documented my learnings in a team wiki and conducted knowledge-sharing sessions. I then architected the solution using Socket.io for signaling and Coturn for TURN services.

**Result:** Delivered a working video chat feature in 2.5 weeks supporting up to 6 concurrent participants. The feature became a key differentiator, contributing to a 30% increase in user engagement. I later created internal training materials that helped the team adopt WebRTC for additional features.

## **6. Tell me about a time when you improved the performance of an application significantly.**

**Situation:** Our React-based dashboard was taking 8-12 seconds to load and rendering was sluggish when displaying large datasets (10K+ rows), causing user complaints and a 25% drop in feature adoption.

**Task:** Optimize the application to achieve sub-2-second load times and smooth rendering for large datasets.

**Action:** I conducted performance profiling using React DevTools and Chrome Lighthouse. I identified issues: unnecessary re-renders, lack of code splitting, and inefficient data rendering. I implemented React.memo for component memoization, virtualized lists using react-window for rendering only visible rows, implemented code splitting with dynamic imports, and optimized bundle size by replacing moment.js with date-fns. I also added server-side pagination and implemented Redis caching for frequently accessed data.

**Result:** Load time reduced from 8s to 1.2s (85% improvement), bundle size decreased by 60%, and the dashboard could smoothly handle 50K+ rows. User satisfaction scores increased by 40%, and feature adoption recovered to previous levels within one month.

## **7. Describe a time when you had to handle conflicting priorities from multiple stakeholders.**

**Situation:** During a sprint, the product team wanted a new payment gateway integration, the DevOps team needed critical security patches applied, and the support team was escalating a data export bug affecting 20% of users.

**Task:** As the senior full stack developer, I needed to prioritize these competing demands while maintaining team velocity and stakeholder relationships.

**Action:** I called an emergency meeting with all stakeholders to assess impact and urgency. I created a risk matrix evaluating business impact, technical debt, and user experience. I proposed: immediately fix the data export bug (2 days), apply security patches in parallel (1 day), and defer payment integration to next sprint with a detailed technical spec prepared. I communicated the

rationale clearly, showing how the data bug affected revenue and the security patches were compliance-critical.

**Result:** All stakeholders agreed to the prioritization. We resolved the bug affecting 5,000 users, applied security patches meeting compliance deadlines, and delivered the payment integration in the next sprint with better requirements. This approach became our standard framework for handling conflicting priorities.

## **8. Give an example of when you had to refactor legacy code while maintaining backward compatibility.**

**Situation:** Our REST API had grown organically over 5 years with inconsistent response formats, no versioning, and tightly coupled business logic. We needed to modernize it without breaking 50+ client applications in production.

**Task:** Refactor the API architecture to follow RESTful best practices while ensuring zero downtime and backward compatibility for existing clients.

**Action:** I designed an API versioning strategy using URL versioning (/v1/, /v2/). I created adapter layers to transform legacy responses to new formats and vice versa. I implemented the strangler fig pattern, gradually migrating endpoints to the new architecture while maintaining old ones. I wrote comprehensive integration tests covering all existing client use cases and created migration guides for clients. I also set up API analytics to monitor which endpoints were still using v1.

**Result:** Successfully migrated 80% of endpoints to v2 over 4 months with zero production incidents. Response times improved by 35% due to cleaner architecture. We deprecated v1 after 6 months with only 2 clients requiring migration assistance. The new architecture reduced API-related bugs by 60%.

## **9. Tell me about a time when you identified and prevented a potential security vulnerability.**

**Situation:** During a code review, I noticed our file upload feature wasn't validating file types on the server side, only on the client side. This could allow malicious users to upload executable files or scripts.

**Task:** Assess the security risk, implement proper validation, and ensure no existing vulnerabilities were being exploited.

**Action:** I immediately conducted a security audit of the upload functionality and discovered we were also missing file size limits and proper storage isolation. I implemented server-side MIME type validation, file extension whitelisting, virus scanning using ClamAV, size restrictions, and moved uploaded files to a separate S3 bucket with strict access policies. I reviewed server logs to check for any suspicious upload activity and found none. I then created security guidelines for file handling and added automated security checks in our CI/CD pipeline.

**Result:** Prevented a potential remote code execution vulnerability before it was exploited. The security measures blocked 15+ malicious upload attempts in the first month. I presented the findings in a security workshop, leading to a company-wide security audit that identified and fixed 8 additional vulnerabilities.

## **10. Describe a situation where you had to balance technical debt with feature delivery.**

**Situation:** Our codebase had accumulated significant technical debt with test coverage at 40%, duplicated business logic across services, and a fragile deployment process. Meanwhile, the business was pushing for three major features for an upcoming product launch in 2 months.

**Task:** Deliver the required features while addressing critical technical debt to prevent future velocity slowdown.

**Action:** I quantified the technical debt impact by tracking bug frequency, deployment failures, and development time for recent features. I presented data to leadership showing that our velocity had decreased 30% over 6 months due to technical debt. I proposed a hybrid approach: allocate 70% of sprint capacity to features and 30% to technical debt. I prioritized debt that directly impacted the new features—refactoring shared business logic into a common library, improving CI/CD reliability, and increasing test coverage for critical paths to 80%.

**Result:** Delivered all three features on time for the product launch. Technical debt reduction improved deployment success rate from 75% to 95% and reduced bug reports by 40%.

Development velocity increased by 25% in subsequent sprints. Leadership adopted the 70/30 split as a permanent policy.

