DocuPanda Disaster Recovery and Business Continuity Plan

1. Introduction

1.1 Purpose of the Disaster Recovery Plan The purpose of this Disaster Recovery Plan is to ensure **business continuity** and **data integrity** for DocuPanda in the event of a disaster. This plan outlines the strategies and procedures required to recover critical systems and services, ensuring compliance with regulatory requirements and maintaining service commitments to clients. It focuses on minimizing downtime and protecting long-term stored data, which is essential for DocuPanda's enterprise clients.

1.2 Scope This plan covers disaster recovery procedures for the following critical DocuPanda components:

- **Long-Term Storage Buckets (S3 or Azure):** Ensuring availability and data recovery for long-term document storage.
- **DocumentDB:** Safeguarding metadata and analysis results.

Temporary outages or capacity issues related to document processing servers or delays in ingesting new documents are out of scope**. These issues are managed under DocuPanda's SLA agreements, which include compensation for downtime. Document processing server outages are considered operational rather than disasters requiring recovery under this plan.

This plan addresses both regional and full-system outages. It ensures that DocuPanda clients remain unaffected by failures within their designated geographic region (US or EU).

1.3 Key Stakeholders The following table lists the roles responsible for disaster recovery planning and execution at DocuPanda:

Role	Responsibilities
Chief Security Officer (CSO)	Oversees business continuity, communicates with clients and legal authorities.
Chief Technology Officer (CTO)	Leads technical recovery efforts, liaises with external cloud providers.
Disaster Recovery Coordinator	Manages the overall disaster recovery process and timelines.
Infrastructure Lead	Coordinates technical recovery for systems, including databases and storage.
Cloud Infrastructure Specialist	Monitors cloud services, identifies incidents, and provides technical expertise.
Compliance and Security Officer	Ensures recovery efforts meet regulatory and security standards.
Business Continuity Manager	Focuses on minimizing business impact and coordinating cross-team efforts.

External Stakeholders:

- **Cloud Service Providers:** AWS or Azure

Primary End Users:

- **DocuPanda's Enterprise Clients:** Businesses relying on DocuPanda's document processing and long-term storage solutions.

1.4 Notification Tree

This section describes the escalation and notification process for disaster scenarios.

1. **Incident Detection and Initial Assessment**
   - **Cloud Infrastructure Specialist** identifies and assesses the incident severity.
   - If the incident is critical, the specialist notifies the **Infrastructure Lead**.
2. **Internal Escalation**
   - **Infrastructure Lead** coordinates recovery efforts with the **Disaster Recovery Coordinator**.
   - The **Coordinator** escalates to the **CTO** for strategic oversight.
3. **Business Impact and Compliance**
   - **CTO** consults with the **Business Continuity Manager** to minimize client impact and ensure alignment with business goals.
   - **Compliance and Security Officer** ensures adherence to regulatory requirements and data security protocols.
4. **Client and External Stakeholder Notification**
   - The **CSO** informs affected clients about the incident, including its impact and resolution timeline.
   - The **CTO** engages **Cloud Service Providers** (AWS or Azure) if cloud recovery is necessary.

Visual Notification Tree

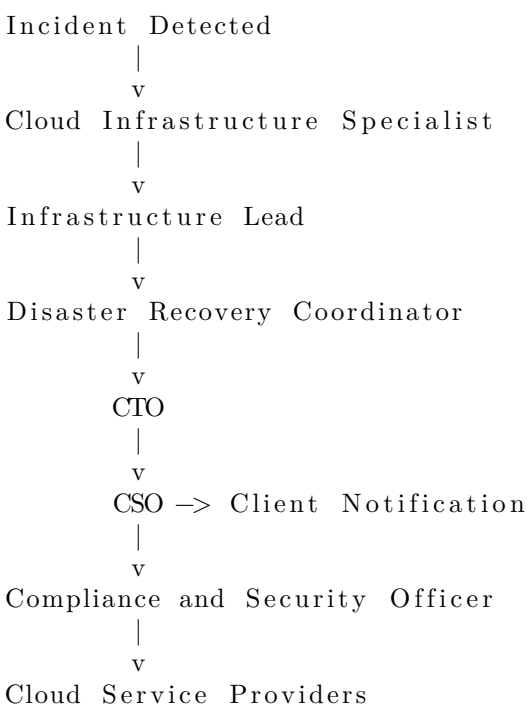

Added Roles Justification:

- **Disaster Recovery Coordinator:** To streamline the recovery process and ensure proper execution.
- **Compliance and Security Officer:** To ensure all legal and regulatory requirements are met during recovery.
- **Business Continuity Manager:** To oversee and manage business impacts, focusing on client-facing services.

---

2. System Overview

2.1 Components Overview

2.1.1 Document Processing Servers DocuPanda's document processing is handled by an **autoscaling group** of stateless servers. These servers dynamically scale based on utilization and perform tasks such as:

- **OCR (Optical Character Recognition)**
- **Parsing**
- **insights generation**: document classification, standardiziation objects, review objects.

Since the servers are stateless, they do not store any data locally. Outages affecting these servers only impact new document ingestion, which is covered under SLA agreements and is not considered a disaster as there is no data loss.

2.1.2 DocumentDB (Metadata and Analysis) DocuPanda utilizes **MongoDB on AWS (DocumentDB)** for metadata storage and analysis. This database stores:

- **Document metadata**
- **Schema-driven insights and analysis results**

The database is configured with **replication** and **sharding** within the client's geographic region. While replication ensures availability, the database does not have cross-region redundancy, adhering to data residency requirements. DocumentDB is a **critical component** as it provides real-time insights and access to previously generated results.

2.1.3 Long-Term Storage Buckets (S3/Azure) Client-uploaded raw files (e.g., PDFs) are stored in region-specific **S3 buckets** or Azure equivalents. Key attributes of the long-term storage:

- **File Types:** Raw documents uploaded by clients
- **Storage Duration:** Based on client-specified expiration policies (ranging from daily deletion to indefinite retention)
- **Regions Supported:** North America (NA), European Union (EU), and Australia (AU)
  - Default storage is in the US unless otherwise specified in client agreements.

These buckets do not directly impact the availability of insights but provide essential data storage for regulatory and compliance purposes.

2.2 Data Flow and Dependencies

1. **File Upload:**
   - Clients upload raw files via the **API Gateway** hosted on AWS.
2. **Processing:**
   - Files are routed to the autoscaling group of **Document Processing Servers** on AWS/Azure.
   - Servers extract insights and store:
     - **Metadata and analysis results** in **DocumentDB**.
     - **Raw files** in the appropriate **Long-Term Storage Bucket**.
3. **Storage & Retrieval:**
   - Clients query **DocumentDB** for insights and metadata.
   - Raw files are retrieved from **S3/Azure storage** when needed, although most clients primarily rely on the processed insights.

Dependencies:

- Document Processing Servers rely on the **API Gateway** and **DocumentDB** for storing results.
- DocumentDB relies on its sharding and replication setup within the same region.
- Long-Term Storage Buckets operate independently but align with client-specific geographic agreements.

2.3 Criticality of Each Component

Component	Impact of Outage	Criticality
Document Processing Servers	New document ingestion and processing halt. Covered under SLA, no data loss.	Low
DocumentDB (Metadata/Insights)	Clients lose access to document insights and metadata retrieval.	High
Long-Term Storage Buckets	Raw files unavailable, but insights remain accessible. Secondary for most clients.	Medium

3. Risk Assessment and Threat Analysis

3.1 Natural Disasters DocuPanda's infrastructure is hosted across the following cloud regions:

- **AWS eu-west-1** (Europe)
- **AWS us-east-1** (North America)
- **Azure East US**
- **Azure West Europe**

Each of these regions comes paired with a paired redundant region for disaster recovery, as per cloud provider's discretion to minimize the likelihood of catastrophic failures and permanent data loss.

A single client is only ever localized to a single geo region. These are only iether US or EU at the time of this revision.

Potential natural disaster risks include:

- **Europe (eu-west-1):** Low risk of hurricanes but possible localized flooding or extreme weather events.
- **North America (us-east-1):** Potential risk of hurricanes, particularly in the Atlantic region, which could affect data centers.

While cloud providers have robust disaster recovery and regional failover mechanisms, localized natural disasters may still cause temporary service disruption or increased latency.

3.2 Cybersecurity Threats DocuPanda faces several key cybersecurity threats:

- **Data Breaches:** Unauthorized access to sensitive document data or metadata stored in DocumentDB or S3 buckets.
- **Unauthorized Access:** Attempts to gain access to internal systems or client data through compromised credentials.
- **Ransomware Attacks:** Potential attacks targeting stored files or metadata to encrypt and ransom data.

Mitigation Measures:

- Role-based access control (RBAC) and strict **IAM policies** on AWS and Azure.
- End-to-end encryption for data in transit and at rest.
- Use of **Web Application Firewalls (WAFs)** to protect the API Gateway from malicious requests.
- Regular security audits and penetration testing.

3.3 Hardware Failures As DocuPanda operates entirely on cloud platforms, **hardware failures** are abstracted by AWS and Azure. Key measures in place include:

- **S3 Buckets and DocumentDB:** Both services offer built-in redundancy and replication to mitigate hardware failure risks.
- **Document Processing Servers:** Managed using an **autoscaling group** with integrated **health checks** to ensure failed instances are automatically replaced.
- **Monitoring:** Integration tests periodically check for processing service availability, ensuring rapid detection and mitigation of any hardware or software issues.

3.4 Software Failures Potential software failure risks could include:

- **Deployment Errors:** Misconfigured services or bugs introduced during updates to the processing servers.
- **Data Corruption in DocumentDB:** Although unlikely, bugs in data handling or querying could cause data inconsistencies.
- **API Gateway Failures:** Issues with routing requests, possibly due to configuration or third-party library updates.

Mitigation Measures:

- All deployments follow a **CI/CD pipeline** with automated testing.
- Rollback mechanisms to quickly revert to stable builds in case of failure.
- Frequent backups of DocumentDB to ensure quick recovery in case of corruption.

---

4. Disaster Recovery Objectives

4.1 Recovery Time Objectives (RTOs) The Recovery Time Objective defines the maximum acceptable downtime for each critical component following a disaster. The following RTOs apply to DocuPanda's systems:

Component	RTO	Justification
DocumentDB (Metadata and Analysis)	4 hours	Critical for providing real-time insights and document metadata access.
Long-Term Storage Buckets	12 hours	Raw document storage, secondary to metadata and insights availability.
Document Processing Servers	6 hours	Affects new document processing; impacts are operational, covered by SLA.

4.2 Recovery Point Objectives (RPOs) The Recovery Point Objective defines the maximum acceptable data loss, typically measured as the time between the last successful backup and the disaster. The following RPOs apply:

Component	RPO	Justification
DocumentDB (Metadata and Analysis)	1 hour	Ensures minimal data loss for metadata and insights critical to clients.
Long-Term Storage Buckets	12 hours	Raw document uploads are less frequently accessed; data loss tolerance is higher.

4.3 Service Level Agreements (SLAs) SLAs are handled separately and are **beyond the scope of this document**. They define the uptime guarantees and compensation models for clients in the event of outages or service interruptions.

---

5. Roles and Responsibilities

5.1 Disaster Recovery Team Structure The Disaster Recovery Team at DocuPanda consists of specialized roles that coordinate and execute recovery efforts during a disaster. The structure includes the following roles:

Role	Primary Function
Chief Technology Officer (CTO)	Oversees technical recovery, authorizes major decisions, liaises with cloud providers.
Chief Security Officer (CSO)	Leads communication with clients, ensures compliance and security standards are met.
Disaster Recovery Coordinator	Manages the recovery process, ensures timelines are met, and keeps all stakeholders informed.
Infrastructure Lead	Executes recovery for core infrastructure, including DocumentDB and processing servers.
Cloud Infrastructure Specialist	Handles recovery of cloud-based services, including S3 buckets and Azure storage.
Compliance and Security Officer	Monitors regulatory adherence and data security during recovery operations.
Business Continuity Manager	Coordinates business impact mitigation and ensures continuity of critical operations.

5.2 Responsibilities and Contact Information Each team member has specific responsibilities during disaster recovery. The following table outlines these responsibilities and includes placeholders for contact information:

Role	Responsibilities	Contact
Chief Technology Officer (CTO)	Authorize recovery procedures, coordinate with cloud providers.	**REDACTED**
Chief Security Officer (CSO)	Notify clients, ensure data integrity and compliance.	**REDACTED**
Disaster Recovery Coordinator	Oversee the recovery timeline, facilitate cross-team communication.	**REDACTED**
Infrastructure Lead	Restore DocumentDB and processing server operations.	**REDACTED**
Cloud Infrastructure Specialist	Recover S3 and Azure bucket services, monitor cloud health.	**REDACTED**
Compliance and Security Officer	Validate recovery actions meet regulatory and security standards.	**REDACTED**
Business Continuity Manager	Ensure operational impact is minimized and critical services resume.	**REDACTED**

**** 6. Backup Strategy****

6.1 Document Processing Servers Backups Since Document Processing Servers are stateless, they do not require traditional backups. Instead, the focus is on maintaining configuration integrity and ensuring seamless redeployment in case of failure.

- **Configuration Management:**
  - **Daily Configuration Snapshots:** Configuration files and server environment variables are captured daily and stored in a secure, version-controlled repository.
  - **Immutable Infrastructure:** Server instances are deployed from pre-validated images using an automated CI/CD pipeline.
- **Code and Dependency Management:**
  - **Continuous Integration Backups:** Application code and all third-party dependencies are automatically archived with each successful build.
- **Recovery Strategy:**
  1. **Automatic Remediation:** Failed instances are replaced by autoscaling groups.
  2. **Manual Redeployment:** If automation fails, manually redeploy using the latest stable configuration.
  3. **Post-Recovery Validation:** Conduct integration tests using real client data to ensure service integrity before scaling operations.

6.2 DocumentDB Backups DocumentDB is a critical system for storing metadata and analysis results, requiring rigorous backup procedures.

- **Backup Type:**
  - **Incremental Backups:** Captures data changes every hour.
  - **Full Daily Backups:** A complete snapshot is taken daily to ensure rapid recovery.
- **Storage and Retention Policy:**
  - Backups are retained for **30 days** in primary and secondary (cross-region) locations.
- **Recovery Testing:**
  - Backups are tested **monthly** in a sandbox environment to validate recovery speed and data integrity.

6.3 Long-Term Storage Bucket Backups (Revised)

DocuPanda's long-term storage buckets hold raw documents uploaded by clients. Ensuring their availability and integrity is crucial for compliance and historical record-keeping. This section outlines the backup strategy and integrity verification processes for long-term storage.

6.3.1 Backup Strategy and Replication

- **Real-Time Replication:**
  All data is replicated in real-time across multiple geographic regions (primary and secondary locations) to ensure high availability and redundancy. This provides immediate failover capabilities during regional outages.
- **Backup Frequency:**
  A point-in-time snapshot of the entire bucket is created **daily** to safeguard against accidental deletions, data corruption, or ransomware attacks.

6.3.2 Integrity Checks and Validation Ensuring data integrity across primary and secondary storage locations is a priority. The following processes are implemented:

- **Scheduled Integrity Checks:**
  A system-wide **checksum validation** is performed **weekly** on all replicated files. This process verifies that files in secondary regions match their primary counterparts, ensuring no corruption or data loss during replication.
- **Automated Alerts for Anomalies:**
  If discrepancies are detected during these checks (e.g., mismatched checksums), automated alerts are sent to the **Cloud Infrastructure Specialist** for immediate investigation.
- **Manual Spot Checks:**
  Monthly manual audits are conducted on a random subset of stored files to validate the accuracy of automated checks and ensure overall system reliability.

6.3.3 Recovery Procedures In the event of data loss or corruption in the primary region:

1. **Failover Activation:**
   - The system will automatically redirect requests to the secondary storage location within **15 minutes** of incident detection.
2. **Restoration of Primary Region:**
   - Once the primary region becomes available, the latest replicated data from the secondary region is synced back to ensure full restoration.
3. **Data Verification Post-Restoration:**
   - After restoration, checksum validation is conducted on all restored files to confirm data integrity.

6.3.4 Retention Policies Retention policies are determined based on client requirements:

- **Default Retention Policy:**
  Data is retained for **30 days** in both primary and secondary locations unless a custom policy is specified by the client.
- **Client-Defined Retention:**
  Clients may opt for custom retention policies, such as daily deletion or indefinite retention. These policies are enforced through automated lifecycle rules.

6.3.5 Testing and Validation

- **Quarterly Disaster Simulations:**
  Recovery scenarios for long-term storage are simulated quarterly to test failover mechanisms, verify data consistency, and measure recovery times.
- **Audit Logs:**
  All validation activities, including integrity checks and manual spot audits, are logged and reviewed during **annual compliance audits** to ensure alignment with regulatory requirements.

6.4 Backup Frequency and Retention Policy

Component	Backup Frequency	Retention Period
Document Processing Servers	Configuration Snapshots: Daily	N/A – Stateless Recovery
DocumentDB	Incremental: Hourly	Full Backup Retained for 30 Days
Long-Term Storage Buckets	Immediate replication	Client-defined or 30 days (default)

6.5 Backup Testing and Validation Regular testing and validation are essential to ensure that backups are reliable and meet recovery objectives.

- **Testing Frequency:**
  - Full backup restoration tests for **DocumentDB** are conducted **monthly**.
  - Real-time replication for long-term storage buckets is tested **quarterly**.
- **Validation Process:**
  1. **Integrity Checks:** Ensure that all data is complete and uncorrupted after restoration.
  2. **Performance Tests:** Verify that the restoration process meets RTO and RPO requirements.
  3. **Audit and Logging:** All tests are logged, and results are reviewed during the annual disaster recovery audit.
- **Corrective Actions:**
  If any issues are identified during validation, they are logged as action items, and resolution is prioritized within **30 days**.

6.6 Backup Security

6.6.1 Encryption

- All backups are encrypted using **AES-256** for data at rest.
- Data in transit during replication and backup storage transfers uses **TLS 1.2+**.

Access Control:

- Backup systems employ **role-based access control (RBAC)** to restrict access to authorized personnel only.
- Regular access audits ensure compliance with security policies.

Improvements Made:

1. Added **default retention policies** to reduce dependency on client-defined parameters.
2. Enhanced focus on **testing and validation**, including detailed recovery and validation processes.
3. Included specific measures for **backup security** to meet regulatory and compliance requirements.
4. Introduced a **post-validation corrective action process** to address issues quickly.

---

7. Disaster Recovery Procedures

This section outlines the steps to be followed in the event of a disaster, detailing how each component is recovered and how communication is managed.

7.1 Incident Detection and Initial Response

- **Detection:** Incidents are detected through automated monitoring systems, including health checks and anomaly detection in cloud services.
- **Initial Response:**
  1. **Cloud Infrastructure Specialist** assesses the scope of the incident.
  2. If the incident qualifies as a disaster (e.g., data corruption, regional cloud failure), they escalate to the Chief Technology Officer (CTO)** or designated **Incident Commander**.
  3. Immediate actions are taken to contain the impact, such as redirecting traffic or disabling affected services.

7.2 Activation of the Disaster Recovery Plan

- **Authorization:** The **CTO** or **Incident Commander** decides to activate the Disaster Recovery Plan.
- **Team Mobilization:** All relevant roles are notified, and recovery procedures are initiated.
- **Client Notification:** The **Chief Security Officer (CSO)** informs affected clients about the situation, providing estimated recovery times and ongoing updates.

7.3 Recovery Procedures by Component Each system component has a tailored recovery process to ensure minimal downtime and data loss during a disaster.

7.3.1 Document Processing Servers Recovery **Objective:** Resume document ingestion and processing with minimal disruption.

Recovery Steps:

1. **Instance Removal and Replacement:**
   - Automatically remove unhealthy instances from the autoscaling group.
   - Deploy new instances using pre-configured images.
2. **Configuration Validation:**
   - Load and validate the latest configuration from the version-controlled repository.
   - Ensure proper environment variables are applied.
3. **Integration Testing:**
   - Conduct end-to-end tests to validate the ingestion and processing pipeline.
   - Use a subset of real data to verify correctness.
4. **Gradual Scaling:**
   - Start with a limited number of instances to monitor performance.
   - Gradually scale to full capacity once validated.

7.3.2 DocumentDB Recovery **Objective:** Restore access to document metadata and insights with minimal data loss.

Recovery Steps:

1. **Backup Restoration:**
   - Restore from the latest available **hourly incremental** or **daily full** backup, depending on the extent of the data loss.
2. **Data Integrity Verification:**
   - Run automated checks to verify data consistency against pre-disaster records.
   - Perform manual spot checks on critical data sets.

3. **Service Reconnection:**
   - Re-establish connections between DocumentDB and dependent services (e.g., document processing servers, API gateway).
4. **Performance Testing:**
   - Execute performance tests to ensure restored databases meet response time requirements.

7.3.3 Long-Term Storage Buckets Recovery **Objective:** Ensure the availability of raw client files and compliance with data retention policies.

Recovery Steps:

1. **Cross-Region Failover:**
   - If the primary region is unavailable, reconfigure systems to access replicated data in the secondary region.
2. **Primary Region Restoration:**
   - Once the primary region becomes available, sync any changes from the secondary region to maintain consistency.
3. **File Integrity Verification:**
   - Perform checksum comparisons between primary and secondary copies to ensure data integrity.
4. **Client Notification:**
   - Inform affected clients of the status of their stored data, including any discrepancies resolved during recovery.

7.3.4 System-wide Validation After individual components are restored:

1. **Integrated Testing:**
   Conduct tests simulating typical user workflows to validate overall system functionality.
2. **Client Data Verification:**
   Allow clients to verify the integrity of critical data through their dashboards or API access.
3. **Approval for Full Restoration:**
   Final sign-off by the **CTO** or **Incident Commander** to resume full operations.

Improvements Made:

1. **Detailed recovery steps** for each component.
2. Introduced **system-wide validation** to ensure smooth integration between components post-recovery.
3. Added a focus on **gradual scaling** and **client involvement** in verification processes.

7.5 Client Communication Procedures

In the event of a disaster, timely and transparent communication with clients is critical to maintaining trust and minimizing confusion. The following procedures outline the steps for client notifications during a disaster recovery process.

Notification Timeline

1. **Initial Notification:**
   - **Timeline:** Within **1 hour** of incident classification as a disaster.

- **Sender: Chief Security Officer (CSO)** or designated backup.
- **Purpose:** Inform clients of the incident, its potential impact, and the immediate steps being taken.

2. **Progress Updates:**
   - **Frequency:** Every **2 hours** during the active recovery phase.
   - **Purpose:** Provide ongoing updates, including recovery progress, estimated resolution time, and any changes in impact.
   - **Post-Stabilization Updates:** If the situation stabilizes but is not fully resolved, updates shift to **every 6 hours**.
3. **Post-Recovery Summary:**
   - **Timeline:** Within **24 hours** after full recovery.
   - **Purpose:** Summarize the incident, root cause, recovery process, and preventive measures.

Notification Content and Templates Initial Notification Template:

```
Subject: [Urgent] Service Disruption Notification - [Service Name]

Dear [Client Name],

We are writing to inform you of a [type of incident] currently affecting
 [specific service].

Our team has identified the issue and initiated our disaster recovery procedures.
Below are the key details:

- **Incident Start Time:** [Time/Date]
- **Affected Services:** [List of impacted components]
- **Impact on Your Data/Service:** [e.g., Metadata temporarily inaccessible, no data loss]
- **Estimated Resolution Time:** [Estimate or "Under "Assessment]

We are committed to keeping you updated on our progress.
Our next update will be provided
at approximately [next update time].

We apologize for the inconvenience and appreciate your understanding as we work to resolve
the issue.

Best Regards,
[CSO Name]
Chief Security Officer, DocuPanda
[Contact Information]
```

Progress Update Template:

```
Subject: [Update] Service Recovery Progress - [Service Name]

Dear [Client Name],

This is an update regarding the ongoing [type of incident] affecting
[specific service].
Our recovery efforts are progressing as follows:

- **Current Status:** [e.g., Backup restoration in progress, 50% of services operational]
- **Estimated Time to Resolution:** [Updated estimate]
```

```
- **Impact Updates:** [Changes in impact, if any]

We will continue to provide updates at regular intervals.
Our next update is scheduled for
[next update time].

Thank you for your patience and understanding.

Best Regards,
[CSO Name]
Chief Security Officer, DocuPanda
[Contact Information]
```

Post-Recovery Summary Template:

```
Subject: [Resolved] Service Disruption Update - [Service Name]

Dear [Client Name],

We are pleased to inform you that the service disruption affecting [specific service]
has been fully resolved as of [time/date]. Below is a summary of the incident:

- **Incident Summary:** [Type of incident and affected services]
- **Root Cause:** [E.g., Regional cloud outage, system misconfiguration]
- **Recovery Actions Taken:** [Detailed actions performed]
- **Impact Duration:** [Total downtime experienced]
- **Data Integrity:** [E.g., No data loss, all systems verified]

To prevent similar incidents in the future, we are implementing the following measures:
[Preventive Actions, e.g., enhanced monitoring, additional redundancy]

If you have any questions or concerns, please 'dont hesitate to reach out.

Best Regards,
[CSO Name]
Chief Security Officer, DocuPanda
[Contact Information]
```

Escalation and Client Escalation Paths In critical cases where service recovery exceeds expected timelines or significantly impacts client operations:

- **Dedicated Client Manager Contact:** Clients are provided a direct escalation path to a senior contact, ensuring priority handling of their concerns.

Improvements Made:

1. Introduced **precise timelines** for each communication stage.
2. Provided **predefined templates** to streamline client communication and ensure consistency.
3. Added an **escalation path** for handling high-priority client concerns during extended recovery efforts.

8. Testing and Maintenance

A robust Disaster Recovery Plan (DRP) requires regular testing and updates to ensure its effectiveness during a real disaster. This section outlines the procedures and schedules for disaster recovery testing and plan maintenance.

8.1 Disaster Recovery Drills and Exercises Regular drills are essential to validate the DRP's effectiveness and to prepare the disaster recovery team.

- **Frequency:**
  Full-scale disaster recovery drills are conducted **biannually**. In addition, targeted component tests (e.g., DocumentDB restoration) are performed **quarterly**.
- **Types of Drills:**
  1. **Tabletop Exercises:**
     - Simulated scenarios are discussed with the disaster recovery team to evaluate their responses and decision-making processes.
     - Focus on coordination, communication, and process adherence.
  2. **Simulation Drills:**
     - Specific disaster scenarios (e.g., DocumentDB failure, S3 region outage) are simulated in a controlled environment.
     - Actual failover and recovery procedures are executed to validate recovery timelines and data integrity.
  3. **Unplanned Drills:**
     - Without prior notice to the recovery team, a simulated failure is triggered to assess real-time readiness and response.
- **Evaluation Criteria:**
  - Recovery Time Objective (RTO) and Recovery Point Objective (RPO) adherence.
  - Effectiveness of communication and role execution.
  - Data integrity and system performance post-recovery.

8.2 Review and Update Schedule

- **Annual Plan Review:**
  The DRP is reviewed **annually** or after significant infrastructure changes to ensure alignment with current operations and compliance standards.
- **Post-Drill Updates:**
  Following each disaster recovery drill, a comprehensive evaluation is conducted. Feedback from the drill is used to update the DRP, addressing any identified gaps or inefficiencies.
- **Change Management Integration:**
  Updates to system architecture, third-party services, or operational procedures automatically trigger a review of the DRP to incorporate new dependencies or risks.

8.3 Post-Test Evaluation Effective post-test evaluation is essential to continuously improve the Disaster Recovery Plan (DRP) and ensure its effectiveness during real incidents. This section outlines the steps and processes for evaluating recovery drills and simulations.

8.3.1 Performance Metrics Each test is evaluated against predefined metrics to determine the success of recovery procedures.

- **Recovery Time Objective (RTO):**
  Measure the actual time taken to restore each critical component. Compare against the target RTOs.
- **Recovery Point Objective (RPO):**
  Assess the extent of data loss by comparing restored data with the most recent backup. Ensure alignment with RPO targets.
- **System Availability and Performance:**
  Verify that system performance post-recovery meets operational standards (e.g., response times, throughput).

8.3.2 Incident Reporting and Documentation After each drill or real recovery event:

1. **Incident Report Compilation:**
   Prepare a detailed report that includes:
   - **Scenario Description:** What was simulated or recovered.
   - **Key Actions Taken:** Step-by-step actions during recovery.
   - **Metrics Achieved:** RTO, RPO, and system performance compared to targets.
   - **Challenges and Delays:** Highlight any areas where performance deviated from expectations.
2. **Stakeholder Review:**
   Share the report with relevant stakeholders, including the CTO, CSO, and Business Continuity Manager, for feedback and discussion.

8.3.3 Corrective Action Process Identified gaps or inefficiencies are logged as action items with clear accountability and timelines for resolution.

- **Gap Analysis:**
  Review each challenge or delay to determine root causes.
- **Action Item Assignment:**
  Assign specific owners for each corrective action, ensuring accountability.
- **Resolution Timelines:**
  Set a deadline for each action item, typically within **30 days** of the evaluation.
- **Follow-Up and Tracking:**
  Track progress during regular review meetings. Ensure all actions are completed before the next scheduled drill.

8.3.4 Continuous Improvement Post-test evaluations contribute to an iterative improvement process:

1. **Plan Updates:**
   Revise the DRP based on findings and incorporate new best practices or changes in infrastructure.
2. **Knowledge Sharing:**
   Conduct internal training sessions to ensure the recovery team is familiar with updates and improvements.
3. **Technology and Process Adjustments:**
   Integrate new technologies or refine processes to enhance recovery capabilities.

8.3.5 Audit and Compliance Review

- **Audit Logs:**
  Maintain comprehensive logs of all drills, tests, and evaluations for audit purposes.
- **Regulatory Alignment:**
  Ensure that all updates and corrective actions maintain compliance with relevant regulations (e.g., HIPAA, GDPR).

---

9. Third-Party Coordination

Effective coordination with third-party providers is crucial to ensuring a seamless disaster recovery process. This section outlines DocuPanda's approach to working with its cloud service providers during a disaster scenario.

9.1 Cloud Service Providers (S3/Azure)

DocuPanda leverages AWS and Azure cloud services for critical operations. These providers ensure high availability and scalability of DocuPanda's infrastructure.

Primary Services Used:

- **Amazon Web Services (AWS S3):** For secure and scalable long-term data storage.
- **Microsoft Azure:** For data redundancy and failover support.

Service Agreements and SLAs:

- **Uptime Guarantees:** Both providers offer a **99.99% uptime SLA**.
- **Premium Support Channels:**
  - AWS Premium Support ensures 24/7 incident resolution.
  - Azure Premier Support provides priority handling for service disruptions.

Failover Strategies:

- Multi-region replication of S3 buckets and Azure storage ensures no data loss during outages.
- Automated failover configurations redirect operations to unaffected regions, minimizing downtime.

---

10. Compliance and Legal Considerations

DocuPanda is committed to maintaining full compliance with both **HIPAA** (Health Insurance Portability and Accountability Act) and **GDPR** (General Data Protection Regulation). This ensures the highest standards of data privacy, security, and legal adherence for all protected health information (PHI) and personal data processed within our systems.

10.1 Regulatory Requirements

DocuPanda ensures compliance with **HIPAA** and **GDPR** regulations, safeguarding client data while maintaining legal adherence.

HIPAA Compliance

- **Security Rule:** Administrative, physical, and technical safeguards for PHI, including encryption and access controls.

- **Breach Notification Rule:** Rapid incident response and breach reporting within **60 days**, as required.

GDPR Compliance

- **Data Minimization:** Collect only necessary data, with deletion endpoints for user data.
- **Breach Notification:** Report breaches to the appropriate authorities within **72 hours**.

Encryption Standards:
All data is encrypted using **AES-256** for storage and **TLS 1.2+** for transit, ensuring robust security.

GDPR Compliance GDPR governs the collection, processing, and storage of personal data of EU residents, ensuring their privacy and data rights.

- **Lawfulness, Fairness, and Transparency:**
  Requires that data processing activities be lawful and transparent to users.
  - Implementation: DocuPanda provides clear **Privacy Notices** and obtains explicit **consent** for data processing when required.
- **Data Minimization and Purpose Limitation:**
  Personal data must be collected only for specific, legitimate purposes and retained no longer than necessary.
  - Implementation: DocuPanda ensures that PHI and personal data are retained only as long as necessary for processing, with automated deletion mechanisms.
- **Rights of Data Subjects:**
  GDPR grants individuals the right to access, rectify, delete, or restrict the processing of their personal data.
  - Implementation: DocuPanda provides self-service tools and support channels for users to exercise their rights, including the **Right to be Forgotten** and **Data Portability**.
- **Data Breach Notification:**
  Requires notification to supervisory authorities and affected individuals within **72 hours** of discovering a data breach.
  - Implementation: DocuPanda's **Incident Response Plan** includes GDPR-specific notification protocols.
- **Data Protection Impact Assessments (DPIAs):**
  Required for processing activities that pose high risks to individuals' privacy.
  - Implementation: DPIAs are conducted for all new or significantly changed data processing operations, ensuring risk mitigation and compliance.

10.2 Data Privacy and Security Regulations DocuPanda's security framework is designed to meet or exceed the data privacy and security requirements under both HIPAA and GDPR.

Data Encryption and Key Management

- **HIPAA and GDPR Alignment:**
  All data at rest and in transit is encrypted using **AES-256** and **TLS 1.2+**, ensuring that sensitive information is secure.
  Encryption keys are managed via **AWS Key Management Service (KMS)**, with strict controls over key access and rotation.

Access Control and Identity Management

- **HIPAA:** Access to ePHI is limited through role-based access controls and mandatory multi-factor authentication (MFA).

- **GDPR:** Personal data is accessible only to authorized personnel, with access logs maintained for auditing purposes.

Data Anonymization and Pseudonymization

- **GDPR Compliance:**
  Data is anonymized or pseudonymized where possible to protect user privacy and reduce risks associated with data breaches.

Monitoring and Auditing

- **HIPAA:**
  Comprehensive logs of data access and system activity are retained for **auditability** and to detect unauthorized access.
- **GDPR:**
  Regular **privacy impact assessments** and ongoing monitoring of data flows ensure compliance with GDPR's accountability principle.

International Data Transfers

- **GDPR Compliance:**
  For any data transfers outside the EU, DocuPanda uses approved mechanisms such as **Standard Contractual Clauses (SCCs)** to ensure adequate protection of personal data.

---

11. Appendices

11.1 Glossary of Terms **Disaster Recovery Plan (DRP):** A documented process or set of procedures to recover and protect a business IT infrastructure in the event of a disaster.
Recovery Time Objective (RTO): The maximum acceptable amount of time to restore a system after a disruption.
Recovery Point Objective (RPO): The maximum acceptable amount of data loss measured in time.
Protected Health Information (PHI): Any information about health status, provision of healthcare, or payment for healthcare that can be linked to an individual.
Multi-Factor Authentication (MFA): A security system that requires multiple methods of authentication to verify a user's identity.
DocumentDB: A scalable database service for metadata and analysis in DocuPanda. —

11.2 Checklist for DR Scenarios

Step	Description	Completed (Y/N)
Incident detection	Verify automated alerts and classify the incident.	
Notification	Notify relevant team members and stakeholders.	
DR Plan activation	Obtain approval from the Incident Commander to activate the DR Plan.	
Component recovery	Execute recovery procedures for affected components (servers, databases, etc.).	
Data integrity checks	Validate recovered data against integrity standards.	
System testing	Perform integration tests to ensure full functionality.	
Communication	Provide updates to internal and external stakeholders.	
Post-recovery review	Conduct a review meeting to analyze the incident and update the DR Plan.	

11.3 Document Revisions and Change Log

Version	Date	Author	Changes	Comments
1.0	May 3, 2024	Uri Merhav	Initial draft of the Disaster Recovery Plan	First authored version
1.1	August 1, 2024	Nitai Dean	Minor typo corrections in Section 6.2	Corrected spelling in "DocumentDB"