# Business Continuity Plan - Summary of Policies for CommCare Platform

This document describes how we will handle business continuity and disaster recovery in accordance with our Business Continuity Policy.  This is a summary of these policies that are specific to the CommCare Platform.

## Business Impact Analysis

A business impact analysis (BIA) is used to determine whether a given application, system or component needs to be included in continuity planning.

The Dimagi BIA process will be driven off of the following:

- Business Impact Analysis
- Business Impact_Analysis_Worksheet.pdf

Systems are inventoried through the cloud resource inventory and the asset management system, and each entry there indicates whether the system is business critical. Applications which have significant financial impact, contractual obligations or regulations associated with them are deemed business critical. Business critical systems shall be independently tracked and a plan for each shall be developed that ensures that they meet the BCP requirements and client SLA's.

## Dimagi Internal Systems

The following outlines measures expected to maintain normal operations around day to day Dimagi activities.

### Data Loss

Data for active projects and content shall be stored in systems that are resilient to data loss, such as Google Workspace, Dropbox, Github, JIRA, and Confluence. These systems maintain backups and support a variety of workflows that allow the Dimagi team to continue to work and update data which gets synchronized when the system comes back online.

### Availability

All internal systems shall be inventoried through asset, vendor or application inventories.

Systems that need to be highly available shall be identified through the business impact analysis process described above. The systems that require high availability are those which support production external systems, including AWS management environments. These have been designed to be operable in multiple AWS regions and need to be tested in the primary availability zones that are in use with Dimagi systems.

Third party systems shall be tracked through the Vendor Management process and the availability SLA's reviewed. Where needed, availability SLA's will be considered in choosing appropriate vendors.

# Dimagi Product - CommCare

The following outlines measures expected to ensure high availability for client systems.

## Data Loss

Data shall be stored in data sources that can be backed up. This may be through platform provided capabilities such as AWS EBS and RDS snapshots, leveraging versioned, and distributed storage. Backups shall be tested for completeness on a quarterly basis, and a complete restore shall be tested annually. These exercises include both the primary operation region and the standby region. The results of these tests are evaluated against the Business Continuity metrics and reviewed for possible improvements for future backups and test exercises.

## Disaster Recovery Strategies

The overall disaster recovery (DR) strategy of Dimagi is summarized below and documented in more detail in the supporting sections. These scenarios and strategies are consistent across the technical layers: devops, data, reporting, etc.

### This disaster recovery plan provides:

- Guidelines for determining plan activation;
- Technical response flow and recovery strategy;
- Guidelines for recovery procedures;
- References to key Business Resumption Plans and technical dependencies;
- Rollback procedures that will be implemented to return to standard operating state;
- Checklists outlining considerations for escalation, incident management, and plan activation.

### The specific objectives of this disaster recovery plan are to:

- Immediately mobilize a core group of leaders to assess the technical ramifications of a situation;
- Set technical priorities for the recovery team during the recovery period;

- Minimize the impact of the disruption to the impacted features and business groups;
- Stage the restoration of operations to full processing capabilities;
- Enable rollback operations once the disruption has been resolved if determined appropriate by the recovery team;

### Data Center Disruption (Zone)

- Operate at a degraded service level
- Heighten monitoring and ensure distributed services are operating at the expected consistency and performance level
- Increase capacity in remaining zones and run deployment and testing routines for staging and production systems as needed, updating deployment configurations to add additional resources

RPO: 6 hours | RTO: 3 hours

### Data Center Disruption (Region)

Backups are regularly copied to a separate Region, and a complete failover can be performed if deemed appropriate after discussions with our customers and partners. The process would require a complete deployment of the production infrastructure to the new region, and data restoration from the latest available backups.

RPO: 24 hours | RTO: 12 hours

## Disaster Recovery Procedures

Dimagi DR procedures are broken into three phases:

#### Response Phase: The immediate actions following a significant event.

- On-call personnel paged via Slack or secondary channels
- Decision made around recovery strategies to be taken
- Full recovery team identified
- Gather data and open tickets to fix any outstanding issues

#### Resumption Phase: Activities necessary to resume services after the team has been notified.

- Recovery procedures implemented
- Coordination with other departments executed as needed

#### Restoration Phase: Tasks taken to restore service to previous levels.

- Operations restored

- Rollback procedures implemented

## Response Phase

Dimagi follows the procedure table below when DR needs to be done.

| Step | Lead | Components |
|---|---|---|
| Identify issue, page on-call personnel | OnCall personnel/Managers | Communicate and escalate the issue. Managers and OnCall personnel must identify the priorities for each service. |
| Identify the team members for DR | CTO, or OnCall personnel if CTO is not available | Notify DevOPS, Principal Engineers, Security Engineers, Product Owners, and Division Leads. |
| Establish a conference line for a bridge call | CTO, or OnCall personnel if CTO is not available | Use Slack #server-fires for text communication and set up a Google Meet meeting, falling back to WhatsApp, SMS, and/or Zoom. |
| Communicate the specific recovery roles and determine which recovery strategy will be pursued | CTO, or OnCall personnel if CTO is not available | Coordinate using the established channels and provide constant updates. |

## Resumption Phase

Dimagi utilizes a distributed approach between multiple AWS availability zones in the same region, and maintains regular backups in a separate AWS region. Note that in AWS, per their published documentation, different Availability Zones (AZ) represent different physical data centers and it is generally believed that it is reasonable to provide redundancy across AZ within a Region as a way to achieve a highly performant and highly available architecture. A Regional AWS outage would be a very low probability event. Dimagi has plans to address both an outage at an AZ level and a Region level.

**Zone:** If one zone experiences performance degradation, outage, or failure, only increased monitoring and minimal manual intervention is required to ensure continuation of service, by increasing resources and capacity as needed to maintain previous performance levels. Once the affected zone becomes available again, any change or increase in capacity that was deemed necessary can be reverted.
**Region:** If the primary Region experiences an outage, the production environment can be easily deployed to the secondary region leveraging our infrastructure as code approach, and data can be restored from the available backups. Once the primary region becomes available again, a decision can be made to either keep the production environment in the secondary region, or to switch back to the main region, depending on the duration of the original outage, the amount of new data to be exchanged, and other considerations.

## Data Center Recovery FAILOVER process

### Initiate Failover

**Zone:**

- Change AWS ELB and RDS configurations to remove the affected zone
- Monitor metrics to ensure performance does not drop below acceptable thresholds
- If needed, run deployment routines to increase available resources for affected services

**Region:**

- Deploy the production environment to the secondary region
- Restore data from backups for needed systems (RDS, CouchDB, ES)
- Update AWS ELB and network configurations
- Update infrastructure codebase and monitoring configurations

### Test Failover

- Verify that performance metrics are consistent with expected levels

### Complete Failover

- Create a Retrospective using the designated Form to document all steps taken
- Communicate outstanding issues, fixes, timeline, and steps taken to Management and Support so that customers and partners can be promptly notified and updated

# Restoration Phase

**Zone:** Once the affected zone becomes operative again, all distributed systems will transparently restore traffic to the previously affected nodes and manual configuration changes made to AWS ELB and RDS can be safely reverted. During this restoration process, extra attention should be paid to the metrics to ensure that performance levels remain stable.
**Region:** Once the primary region becomes available again, we can maintain the production environment in the secondary region, or switch it back to the primary region, consequently

adjusting the backup procedures, monitoring configurations, codebase, and available resources as necessary.

## Data Center Recovery FAILBACK process

### Initiate Failback

**Zone:**

- Monitor load balancing and synchronization processes with the recovered nodes
- Remove all additional nodes that are no longer needed
- Monitor metrics to ensure performance does not drop below acceptable thresholds
- Restore previous AWS ELB and RDS configuration

**Region:**

- Maintain the production environment in the current region or synchronize new data with the original region
- Verify and update backups procedures depending on the new active region
- Update/revert AWS ELB and network configurations as needed
- Update/revert infrastructure codebase and monitoring configurations as needed
- Remove or update resources in the new and old region as necessary

### Test Failback

- Verify that performance metrics are consistent with expected levels.

### Complete Failback

- Update the previously created Retrospective to document all of the additional steps taken
- Communicate timeline and steps taken to Management and Support so that customers and partners can be notified