Trice

# Business Continuity Plan

A business continuity plan is a roadmap for continuing operations under adverse conditions. This could be a failure with technology, such as a downed service, or a natural disaster, such as a storm or power outage.

## Background

The Tricefy service is a subscription SaaS service that resides in the cloud (Amazon Web Services). Cloud platforms, such as AWS, provide geographically dispersed data centers around the world. These platforms also provide capabilities that support availability and elasticity that are the cornerstones of the Tricefy Business Continuity plan. The Tricefy architecture was designed to capitalize on these fundamentals from the onset.

## Core Principles

### Availability

Available applications remove single points of failure through redundancy and resilient design. A Business Continuity plan expects each service to potentially go down at different times; the plan needs to specify how the system remains operational and usable when these service interruptions occur. Tricefy remains operational when any of these single failure scenarios occur:

- Server instance not functioning: Tricefy uses multiple application servers in each tier. When a server goes down, the other hosts in the tier will take over the needed processing while a new server spins up. See the Tricefy Security White paper and the Tricefy Architecture presentation for descriptions of the tiers in the Tricefy architecture.
- Network outage (single Availability Zone): Each Tricefy tier runs in multiple Availability Zones (multi-AZ). If there is a network outage in one zone, then all requests and associated processing will occur in the remaining Availability Zones.
- Database offline: The Tricefy data tier includes a multi-AZ RDS Database deployment. Multi-AZ RDS details are described at https://aws.amazon.com/rds/details/multi-az/.
    - When a Multi-AZ DB Instance is provisioned, a primary database instance is created and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically district, independent infrastructure and is engineered to be highly reliable.
- Object Storage inaccessible (single Availability Zone): Objects are redundantly stored on multiple devices across multiple facilities in an AWS S3 region, so this single failure possibility is mitigated.
- Job Queues inaccessible (single availability Zone): Amazon SQS stores all message queues and messages within a single, highly available AWS region with multiple redundant Availability Zones (AZs), so that no single computer, network, or AZ failure can make messages inaccessible.
- Caching service down (single-availability zone): If the caching service is unreachable, the system will switch to a disk-based caching mode backed by S3 leaving the system completely operational.



Trice Imaging, Inc.
1343 Stratford Ct
Del Mar, CA 92014 USA

+ 1 858 344 1531
info@triceimaging.com
www.triceimaging.com

- DNS Resolution failures: IP addresses are used for many service requests within Tricefy to avoid a heavy reliance on DNS servers which may become compromised or inaccessible.

### Scalability

When a system is scalable, it scales horizontally or vertically to manage increases in load while maintaining consistent performance. Cloud systems like Tricefy scale horizontally: horizontal scaling adds more machines (typically of the same size) to the fleet of instances used by the application. This allows Tricefy to be elastic in servicing the throughput needs of end users.

## Catastrophic Scenarios

### Data Corruption

- Database corruption: For business continuity:
  - Take regular database snapshots and persist in multiple regions
  - Verify that database restoration can take place with minimal downtime
- Object storage elements missing:
  - File objects marked for deletion will, by default, be purged after 35 days (coinciding with database daily backup retention).

### Application Failure

- A comprehensive alerting/monitoring system exists for all tiers of the Tricefy system so that there is quick notification of any application failures
- If an emergency application redeploy is necessary:
  - New machine images (AMIs) can quickly be created from the staging environment
  - System redeployment of machine images for the production environment is fast and easy

## Disaster Planning Strategies

1. Conduct a risk assessment for the system (as a whole), for each software service used and for each application tier developed
2. Choose the optimal approach for high availability, while balancing cost, complexity, and risk
3. Design for failure, starting with the application architecture
4. Employ rigorous software development processes to ensure that high availability requirements are implemented and verified for all software modules and services
5. Implement disaster recovery plans and processes including identifying the owners of each piece of the pipeline
6. Document and test all related processes so they are easily repeatable and verifiable



Trice Imaging, Inc.
1343 Stratford Ct
Del Mar, CA 92014 USA

+ 1 858 344 1531
info@triceimaging.com
www.triceimaging.com

# Teams & Responsibilities

The Business Continuity Management Team is composed of senior technology management and quality assurance management. This team is responsible for:

- The overall plan including specification, architecture, development and validation
- Disaster recovery simulation
- All communications (up and down) regarding any crisis or disaster situation

The Senior Technology Management Team includes the Head of Engineering, Director of Quality & Technical Operations, Product Manager, and the Director of Quality. When a crisis or disaster has been identified, this team will do all the necessary communication and recovery planning.



Trice Imaging, Inc.
1343 Stratford Ct
Del Mar, CA 92014 USA

+ 1 858 344 1531
info@triceimaging.com
www.triceimaging.com