Alabama Department of Labor Unemployment Insurance Program Disaster Recovery Plan White Paper

Introduction

Disaster recovery involves a set of policies and procedures to ***enable the recovery or continuation of vital systems following a natural or human-induced disaster***. Disaster recovery focuses on the technology systems supporting critical business functions. The robustness of a system's Disaster Recovery capability is measured by

- **The time it takes to restore a system** (the recovery time objective, or RTO), and
- **The amount of time over which data could be lost** (the recovery point objective, or RPO)

as a result of the disruption. In both cases, shorter time periods are better.

The RTO is the targeted duration of time within which a business process must be restored after a disaster (or disruption) in order to avoid unacceptable consequences associated with a break in business continuity. It answers the question "How long can I be down for?" The RTO is expressed as an amount of time from the instant at which the failure occurs, and can be specified in seconds, minutes, hours, or days. RTO includes the time spent trying to fix the problem without a recovery, the recovery itself, testing, and the communication to the users.

The RPO is a measure of the maximum time period in which data might be lost if there is a major incident affecting the service. It answers the question, "How much data can I afford to lose?" The RPO is expressed backward in time (that is, into the past) from the instant at which the failure occurs, and can be specified in seconds, minutes, hours, or days. The RPO capability of a system is tied to the times backups are sent offsite: during a disaster recovery, the assumption is that the system will lose all data since its most recent backup. So, for example, if the RPO is one hour, backups must be made and sent off-site at least once per hour.

Objectives

Utilizing the strategies identified below, the Alabama Department of Labor is currently meeting the following objectives for Unemployment Insurance-related data:

- Recovery Time Objective (RTO): 60 minutes
- Recovery Point Objective (RPO): 10 minutes

ADOL Disaster Recovery Strategies

The Alabama Department of Labor stores all critical Unemployment Insurance-related data in the Microsoft Azure cloud. Many of the Department's strategies for Disaster Recovery involve utilizing the services that are built into the Azure cloud:

Regional Datacenters

The Microsoft Azure platform allows ADOL to run cloud services such as databases, caching, and websites in government-specific datacenters located around the country. At any time, one datacenter is designated the primary center for providing services, but should that datacenter experience problems, user traffic will be distributed to service endpoints in different datacenters. Microsoft's Azure Traffic Manager allows ADOL to direct client requests to the most appropriate endpoint based on the health of the endpoints. Traffic Manager is resilient to failure, including the failure of an entire Azure region. Traffic Manager delivers high availability for ADOL applications by monitoring the endpoints and providing automatic failover if an endpoint goes down.

Geo-Replication

All services are deployed to multiple regional datacenters. This includes applications, databases, dependent services, and everything necessary for the operation of ADOL's back-end Unemployment Insurance service.

Geo-Replication enables ADOL to configure a readable secondary database in a different regional datacenter. This secondary database is available for read-only querying/reporting and for failover in the case of a data center outage or the inability to connect to the primary database. If for any reason the primary database fails, or simply needs to be taken offline, ADOL can failover to the secondary databases with no more than 5 seconds of data loss.

Point-In-Time Database Restore

Point-In-Time Database Restore uses automated backups in conjunction with transaction logging to recover a copy of a database to a known good point in time. After the database is restored, ADOL can either replace the original database with the restored database or copy any needed data from the restored data into the original database.

Offline Data Backups

Each night, data is backed up from the cloud and delivered to a location on-site at ADOL.

Disaster Recovery Scenarios

Scenario: Network Outage

Scenario: Parts of the Azure network are inaccessible; application(s) or databases(s) are unavailable.

Mitigation: If one or more application instances become unavailable due to network issues, Traffic Manager will redirect users to available instances of the application – even failing over to a different datacenter.

If an application can't access its data because of an Azure network outage, utilizing database geo-replication and Traffic Manager, the system can automatically failover to an alternate region. Recovery Objectives will be met.

Scenario: Region-wide Service Disruption

Scenario: There is a service disruption of an entire Azure region.

Mitigation: If Microsoft declares a region lost, Azure automatically remaps all of the DNS entries to a geo-replicated region. Recovery Objectives will be met.

Scenario: Azure-wide Service Disruption

Scenario: All Azure regions experience simultaneous disruption.

Mitigation: Widespread service disruptions that span regions should be much rarer than isolated service disruptions that involve dependent services or single regions. However, an Azure-wide service disruption incurs the risk of temporary downtime. Recovery Objectives will not be met.

In case of such an event, ADOL would have to evaluate the situation in consultation with its vendors. ADOL may decide to simply endure the temporary down time, or ADOL may decide to restore data from on premise location, along with applications and dependent services to either an alternative cloud service or an on-premises solution.