HIPAA Disaster Recovery Guide

ATLANTIC.NET

Secure Cloud Services
Managed & Compliant Infrastructure

888-618-DATA (3282)
sales@atlantic.net
www.atlantic.net

Table of Contents

This whitepaper details the key personnel, processes, and procedures required for the preparation and planning of a disaster recovery (DR) scenario. DR is achieved by providing healthcare clients with a resilient and HIPAA-compliant cloud infrastructure service.

This service utilizes a leading disaster recovery technology platform, providing the healthcare client with compute, network, storage, and an application stack which is seamlessly replicated to a target data center location using encrypted transfer techniques.

This technology allows for some of the best-in-class Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO). Data is replicated in near real-time, and the cloud service enables us to utilize an on-demand compute infrastructure to rebuild production environments within a disparate cloud data center location.

Purpose

This document aims to serve as a client reference point; it is an overview of the disaster recovery plan (DRP) and the processes established to ensure a smooth transition into DR operations. We will also discuss the planning of failing back of infrastructure services to source production location.

The aim of the Disaster Recovery Plan (DRP) is to explain:

- What is a disaster recovery scenario?
- When to declare a disaster?
- How does an organization invoke disaster recovery?
- How does communication flow during a disaster scenario?
- What are the key roles and responsibilities of personnel assigned to the recovery team?
- How does the business keep running in the event of a disaster?
- Recovery time objectives – how long can the business operate without critical IT systems?
- Recovery point objectives – to what point in the processing cycle can the provider recover?
- What are the post DR activities? (root cause analysis)
- How can the organization continuously test and evolve the DRP?

The Business Associate Agreement

HIPAA compliance demands that the cloud provider have a full understanding of the client's cloud hosting requirements and offer specific guarantees of compliance to the physical, technical, and administrative safeguards required for HIPAA certification.

Before any disaster recovery planning, work is completed to understand our clients' hosted infrastructure and cloud infrastructure requirements for disaster recovery, as well as the what important ePHI protected data is in scope and must be protected.

It is important to work with health-

care clients to understand:

- What electronic patient health information (ePHI) data is in scope. and the servers, databases, and applications associated with that information
- What is the source computing infrastructure which requires restoration to the recovery location during a DR scenario (this is typically referred to as in scope production servers)

The BAA outlines the shared responsibilities for data protection under HIPAA between the cloud service provider and the client, within which the disaster recovery service can be designed. It enables the creation of service processes, procedures and automation blueprints for recovering services. It enables the provider to design the failover and failback process, and also helps with the agreement of service levels.

Disaster recovery (DR) infrastructure as a service (IaaS) is usually only applicable to client production servers. Test and development infrastructure is typically out of the scope of most DR offerings.

Recovery Time Objective (RTO)

RTO is the maximum tolerable time taken to recover customer systems after a disaster scenario is declared. RTO is measured from the time the hosting services are declared to have become unavailable until the time such services become available and operational per applicable service level agreements.

Recovery Point Objective (RPO)

RPO is the measure of the maximum acceptable data loss recorded by time. It is the maximum allowed age of data files when recovering the service. RPO is measured from the time the first transaction is lost, or from the time the hosting services became unavailable to the time services returned to operational.

Hosted Infrastructure Redundancy Protection

The cloud provider builds and maintains a minimum of N+1 redundancy in their data center technology. This ensures resiliency in system availability in the event of a component failure. N+1 means components (N) have twice the amount required for continued operation. This helps ensure the highest levels of reliability and availability in primary and recovery data centers if any components were to fail.

Regular data center checks are completed on top of the automated monitoring which is in place to ensure the maximum redundancy of the hardware.

Key areas where N+1 is employed include:

- Networks – dual feed power supply (PSU) and failover paired networking stack
- Storage – varying RAID levels to ensure disk redundancy, dual feed PSU and multisite replication technology
- Storage which is highly redundant across multiple regions
- Servers – dual feed PSU, clustering, high availability for critical applications (where applicable)
- Real-time monitoring
- Network load balancing on multiple layers for separation of web frontends, applications, and databases

Even despite all these redundancy protection systems, an outage may still happen. Consideration must be given to a total loss of the primary site, such as a lightning strike, tornado, or flood.

The Disaster Recovery Plan

How to Activate the Disaster Recovery Plan

Cloud providers often request that the client appoints a DR coordinator for all communications relating to the disaster recovery process. It is also advisable to have backup con-

tact in case the usual person is on vacation or unwell.

Invocation of the disaster recovery plan must be agreed upon by one of the pre-approved authorized personnel who must contact the cloud provider using the emergency contact details provided in the DRP.

The Disaster Recovery Process

Once the disaster recovery plan (DRP) has been declared, the case will be escalated at the highest priority to the disaster recovery lead (DRL). The DRL will start the process of activating DR and declaring an official state of disaster. It is during this phase that the DRL must ensure, if applicable, that anyone that is in the primary location at the time of the disaster has been accounted for and evacuated to safety.

An all-hands call will be established by the DR Lead with the aim of:

- Confirming to all that a disaster has occurred
- Explain what has happened and (if known) the nature of the disaster
- Understanding the extent of the disaster
- Understanding the impact the DR scenario has had on production services
- If possible, estimate the expected time it will take to recover from the disaster
- Document the steps taken prior too, and since declaring a disaster

Establish a Command Center

If applicable, the provider must establish a business continuity office separate from the primary data center location. Typically, remote workers will be contacted and may be asked to attend the command center. If the primary production location is not physically affected, then operations can continue from the usual headquarters.

The command center is the focal point for all recovery actions and operations. The command center will also be where all communications to

key stakeholders will be provided. The cloud provider is responsible for keeping the client's senior management team and all command center personnel informed on all available communications channels during the DR.

Establish a Communication Plan

As part of the HIPAA compliance procedures, up-to-date call trees will be used to contact all relevant support teams, client contacts, and important personnel.

The DR lead will need to communicate with various parties to inform them of the disaster and the impact on the business. This may include contacting employees, customers, third-party vendors, and suppliers.

Depending on the DR scenario, it will usually be the technical teams who will receive the first alerts or communications about a disaster. These teams will then contact the required escalation persons using a detailed escalation procedure.

Modern communication methods can be used, including automated text alert systems for critical system outages; employees may also have a company mobile with access to employee contact details, email addresses, and emergency procedures.

To cover all scenarios, the DRP emergency procedures, the DRP escalation path, and the emergency contact details for employees, customers, and third party providers should be stored on physical media in two separate physical locations within fire safes. All DRP procedures must be regularly reviewed and updated accordingly in all locations.

Recovering Servers and Services

After the disaster has been declared and approval has been granted to invoke the DRP, the technical teams will then begin to recover client servers and services. The replication

technology used by the cloud provider leverages vast reserves of computation within the recovery location.

This process might be system-automated, triggering when certain pre-defined disaster conditions are met. The disaster recovery solution utilized by the cloud provider varies as per provider but is typically a source-to-target replication solution.

Technical Teams Recover Services

Technical operational playbooks and disaster recovery procedures are a vital part of HIPAA compliance. When in a DR scenario, these documents contain the building blocks to recover critical business services. Despite recovery often being an automated process, it is still vital to have a step-by-step recovery guide available if issues arise.

The technical playbooks give a detailed step-by-step process for the technical teams to follow. This is a comprehensive technical guide explaining how to restore service into the DR location.

To summarize, this process may contain the following steps:

- Wait for approval from the DR team lead to activate the disaster recovery plan
- All teams will report into the command center via the assigned bridge number and receive recovery activities pertinent to their teams
- The technical infrastructure team will be informed which client services are affected by the outage and are being moved to their alternate DR site
- The operations team will identify the appropriate client credentials and validate that infrastructure systems are being failed over too
  - They must log in to the designated DR infrastructure and manually restore the required infrastructure servers
  - They must monitor the designated client services and manually restore the required client services

- The operations teams will bring up all shared systems and validate that all services are running and operational
- The operations teams will confirm monitoring tools are set up for all of the shared servers/applications and networks
- The operations team will confirm backup scripts are in place and backups are scheduled
- The operations network team will cut over the client DR network tunnels to the alternate DR site to "switch" client secured VPNs to the recovery (DR) location
- The operations network team will modify any DNS entries and change internet routing to send all infrastructure related traffic over to the alternate DR failover site
- Designated system administrators will work with the operations teams to validate all connectivity to all application services and that they are operating as expected
- Validate that all ePHI data is now available in the DR location and available within RTO and RPO service levels

The disaster recovery procedures should contain information about:

- Details of each client server instance types, operating systems, and applications
- List of critical servers with ePHI
- The selected regional failover site for the client's primary site
- Documented VPC subnets and replication network information
- The startup order of the infrastructure services if manual restoration is needed

Monitoring the Recovery Progress

The disaster recovery lead will track the progress, report into the command center and make the necessary notifications to the appropriate teams while in a disaster recovery situation. Updates will be provided using the communication plan detailed above.

The technical teams will track the progress of the recovery using any monitoring tools available. These will monitor the health of the applica-

tions, network, and infrastructure. SNMP monitoring agents are typically already deployed to each system so base level monitoring of the server uptime, operating system information, database information can be checked and cross-referenced with playbooks and procedures.

Monitoring solutions vary from each cloud provider. Monitoring dashboards can provide a visual representation of IT systems that are down, pending, or in a failover state.

Testing

After the service has been recovered into the DR site using the cloud provider's replication technology, the provider must work with its clients to test essential infrastructure and IT services to ensure data consistency and that production services are in line with RTO and RPO objectives.

Arguably the most important step here is the data integrity checks – these ensure that all electronic patient health information (ePHI) is available and protected and that the healthcare organization can access the data and associated applications from the recovery location.

Testing is a two-way process where the client and the provider will need to work together to ensure the service is fully restored.

Root Cause Analysis (RCA)

During a disaster scenario, the recovery team's objectives will always be to implement the disaster recovery plan and recover production services as quickly as possible. Post-scenario, the questions of what happened and why can be asked.

Within seven days of the DR scenario, all of the participants involved, including management, DR leads, technical teams, etc. must meet to discuss what happened. These meetings aim to create a root cause analysis (RCA) report and get a full understanding of the disaster incident from start to finish.

It is crucial this meeting takes place

as soon as possible to ensure the facts are fresh in everyone's minds, but this also serves as a critical factor of HIPAA compliance, as the lessons learned during the DR incident need to be translated into updating the process and procedures followed.

The goals of the meetings are to discover and document:

- Start date & time of the incident
- End date & time of the incident
- Details of the support teams involved
- Description of initial details which caused the decision to invoke DR
- Symptoms of incident
- Impact of incident
- Full analysis of the DR scenario
- Discuss the causes of the incident
- Discuss the resolution details
- What worked well?
- What could be done better?
- What lessons have been learned?
- Can any preventative measures be introduced?
- Discuss the resolution details of how the incident was resolved
- List timings of milestones during incident
- Create the root cause analysis report
- Update and make any changes that may be required to the disaster recovery plan.

Plan Service Relocation and/or Return to BAU

Dependent on the disaster scenario, there may be a requirement to keep services running in the DR site for a lengthy period. During this time, it is essential to carefully plan a strategy for returning service to the primary production site/region. Data in the primary site will most likely be out of sync, so consideration must be made to ensure a complete resync of data so a smooth failback of recovery can take place.

Once the primary site has been brought back online, the technical teams will identify the primary site and start to replicate data back from the recovery site to the primary site. Once data integrity is available, the

replication technology will allow you to failback services from the recovery (DR) site back to the primary site. The failback process is usually an automatic job that runs once invoked by the technical teams.

Again, the failback will be monitored and thoroughly tested upon successful failback. The final step is to ensure that the replication from primary to failover site is reactivated and in sync.

Get Help with HIPAA Compliance

HIPAA Compliant Hosting by Atlantic.Net is SOC 2 & SOC 3 certified and HIPAA & HITECH audited, designed to secure and protect critical healthcare data and records. Get a free consultation today! Call 888-618-3282 or review our solutions at https://www.atlantic.net/hipaa-compliant-hosting/.