PAGELY®

We help Big Brands, Scale WordPress.

WORDPRESS HOSTING MANAGED BY PROFESSIONALS

PAGELY, INC
pagely.com

DISASTER RECOVERY BRIEF

# PAGELY DR AND RECOVERY

DETAILED REGIONAL DISASTER RECOVERY AND ZONE FAILURE SCENARIOS

## QUICK REFERENCE

| Plan Type | Recovery Type | Zone RPO | Zone RTO | Zone/ Instance Failure Covered | Region Failure Coverage | Min # of nodes | Private RDS Required |
|---|---|---|---|---|---|---|---|
| **VPS/ENT Single** | Rebuild | 1-24hr | 1-4hr | Yes | RPO: **1-36hrs** RTO: **1-6hrs** | 1 | No |
| **VPS/ENT + HA** | Failover | 1-15min | 1-5min | Yes | RPO: **1-36hrs** RTO: **1-6hrs** | 2 | No |
| **Cluster with DR footprint** | Failover | 1-15min | <1hr | Yes | Enhanced | 3 | 1 each region |

**RPO** = Recovery Point Objective: The point in time for your data set for recovery

**RTO** = Recovery Time Objective: The amount of time expected to recover from a failure

## Detailed Brief

The DR plans apply to loss of a major Pagely system or Amazon system such as the loss of an EC2 availability zone or region.

We provide multi-AZ HA VPS solutions for workloads that need fault tolerance from a single instance or zone failure.

Every VPS plan has a DR protocol for Pagely to employ when the action is justified. RTO and RPO targets can differ based on the chosen hosting plan. We offer enhanced DR options for an additional cost to meet more strict RTO and RPO goals than what the standard strategies deliver.

In our experience the most frequent types of failures on AWS are not widespread zone or regional outages (although those happen too). The most common scenario is unexpected hardware failure or scheduled maintenance impacting a single EC2 instance. For these cases, our multi-AZ HA solutions offer a substantial advantage in resiliency from the most frequently occurring class of outages.

In the case of a regional failure, invocation of the DR plan is ultimately dependent on the nature and severity of the outage affecting the primary site, as such an event causes a divergence of data into two distinct sets and later reconciliation is necessary for a full recovery. For instance, full recovery time going back to the primary site may actually be lengthened if a DR cutover is performed.

**It's important to agree in advance on an amount of time to elapse before executing the DR steps. At the very least, some manual data merging would become necessary to ensure nothing gets left behind while transitioning a service back home.**

For these reasons, it's important to evaluate the costs and benefits when considering adding additional degrees of redundancy to your web properties.

# FAILURE SCENARIOS AND PLAN TYPES

## Failure scenario: Availability Zone Outage

RDS is multi-AZ, with automatic failover from AWS, no restoration of DB necessary

With either active/active or active/passive options, object cache and session storage is handled by a redis instance on the primary server. In the event of an extended zone outage, that will be repointed by Pagely to the server that is still online.

With either active/active or active/passive options, for extended zone failures, the server that is still online becomes the new primary node and another secondary is built in another available zone and added to the HA DNS recordset.

## VPS-1 / any standard single node option

Server rebuild time ~ 20-35 minutes
File restore time ~ 1-3 HRS, depending on amount of data

## VPS-1 and passive standby node

- Traffic automatically sent to standby by Route53 health checks failing on primary server
- Standby node is already built and running
- Has fresh file replication
- For short term outages, the standby server is kept as is and accepts all traffic. After recovery of primary node, a backwards file sync is performed.
- For a major zone outage, the standby is resized to the same size as the original primary server. The resize event involves a stop/start of the instance, about 2-5 minutes.
- For truly catastrophic zone failures, the server that is still online becomes the new primary node and another secondary is built in another available zone and added to the HA DNS recordset.

## VPS-1+ / any standard HA option

- Traffic is automatically sent to surviving node only by Route53 health checks failing on lost server
- RDS is multi-AZ, with automatic failover from AWS.
- Failover is less troublesome than active/passive, but running at half capacity with one node down.
- For short term outages, the surviving server is kept as is and accepts all traffic. After recovery of primary node, a backwards file sync is performed.
- For a major zone outage, or if the diminished capacity is insufficient, the surviving node is resized to the same aggregate size of the HA pair as normal.

# FAILURE SCENARIO: REGION OUTAGE (LESS COMMON)

## Any type of VPS (single node, a/a, a/p HA)

If there is a failure of the home region and there is no existing DR footprint, Pagely would be building an entirely new server(s) in another region that is operational and begin restoring file and database backups.

Server rebuild times are extended and may be further complicated if there is insufficient instance availability. If a region is down, other AWS users are likely to be building new servers in other regions at the same time, causing a run on instance availability.

For added protection from this scenario, Pagely recommends having customer provide a secondary S3 bucket in an alternate region in case the outage is also impacting S3 in the home region. Our automated backup system will upload a copy of files and database to that bucket as well.

Recovery Objectives for a regional outage without pre-existing DR footprint:

- RPO: **1-36hrs**
  RTO: **1-6hrs**

*Cross Region Replication of Pagely managed S3 backup buckets coming soon

## Cluster with DR footprint

In the case of a regional failure, invocation of the DR plan is ultimately dependent on the nature and severity of the outage affecting the primary site, as such an event causes a divergence of data into two distinct sets and later reconciliation is necessary for a full recovery. For instance, full recovery time going back to the primary site may actually be lengthened if a DR cutover is performed.

**It's important to agree in advance on an amount of time to elapse before executing the DR steps. At the very least, some manual data merging would become necessary to ensure nothing gets left behind while transitioning a service back home.**

Recovery Objectives for a regional outage with DR footprint:

- RPO: **1-15 min**
  RTO: **15min-1hr**
- Continuous replication of DB and application code, content to alternate region enables quicker recovery time.
- Cutover to DR footprint for web traffic is automatic with Route53
- Promotion and scale-up of RDS would be manual
- Scale-up of application servers would be manual
- For extended outages, the DR footprint acts as a seed to build a whole new primary site with original specs of the original site.