

AWS Elastic Disaster Recovery (DRS)

Proof of Concept Checklist and User Guide



Checklist

Section	Status	Notes
Analyze your environment		
Define Workloads for DR		
Review Pre-requisites <ol style="list-style-type: none">1. Network requirements2. Ensure that the source server operating system is supported by AWS3. Linux Installation Requirements4. Windows Installation Requirements		
Define RPO and RTO requirements		
Admin Level Access to the AWS Account		
AWS Account selection for DR Site		
AWS Regions Selection for DR site		
Create dedicated Staging VPC and Subnet in the DR account		
Initialize DRS using AWS DRS Console		
Ensure the on-prem and AWS networking requirements are met		
Install Replication Agents on each source server		
Configure Launch Settings for EC2 servers in DR account		
Perform a Test/Drill Failover		
Validate the Failover is successful		
Test the workload is up and running on AWS		

2. To use AWS Elastic Disaster Recovery (DRS), you first need to set it up in each AWS Region in which you want to use it. (The target region is the same for replication and failover.)

Note: You need to be the admin user of the AWS account, or have a role with the “AWSElasticDisasterRecoveryConsoleFullAccess” permission to set up the service.

3. AWS Elastic Disaster Recovery setup

- a. Before setting up Elastic Disaster Recovery, create a subnet which will be used by Elastic Disaster Recovery as a staging area for data replicated from your source servers to AWS. You must specify this subnet in the Replication Settings template. You can override this subnet for specific source servers in the replication settings. While you can use an existing subnet in your AWS account, the best practice is to create a new dedicated subnet for this purpose.

Note: Best practice is to have a dedicated VPC for staging, failover, and testing.

- b. Set default replication settings. This includes setting up replication servers, volumes and security groups, and additional settings. Then you will review and create the default replication settings. For more information you can go to: <https://docs.aws.amazon.com/drs/latest/userguide/quick-start-guide-gs.html#first-time-setup-gs>

- c. [Add source servers](#)

- i. Networking Requirements: Prior to adding your source servers, ensure that you meet all of the [Network requirements](#). To prepare your network for running Elastic Disaster Recovery, you need to set the following connectivity settings:
 1. [Establish communication over TCP port 443](#)
 2. [Establish communication between the source servers and Elastic Disaster Recovery over TCP port 443](#). Each source server that is added to Elastic Disaster Recovery must continuously communicate with Elastic Disaster Recovery (DRS.<region>.amazonaws.com) over TCP port 443.
 3. [Establish communication between the Staging Area Subnet and Elastic Disaster Recovery over TCP port 443](#)
 4. [Communication between the source servers and the Staging Area Subnet over TCP port 1500](#). TCP port 1500 is needed for the transfer of replicated data from the source servers to the Staging Area Subnet.

- ii. Installation Requirements – Before installing the AWS Replication Agent on your source servers, check that they meet the following requirements:
 - 1. [General Requirements](#)
 - a. [Ensure that the source server operating system is supported by AWS](#)
 - b. Ensure that your setup meets all replication networking requirements as mentioned above.
 - c. Ensure stability of the MAC address. Confirm that the MAC addresses of the source servers do not change upon a reboot or any other common changes in your network environment. Elastic Disaster Recovery calculates the unique ID of the source server from the MAC address. When a MAC address changes, Elastic Disaster Recovery is no longer able to correctly identify the source server. Consequently, replication will stop. If this happens, you will need to re-install the AWS Replication Agent and start replication from the beginning.
 - 2. Source Server Requirements – The following are universal requirements for both Linux and Windows source servers:
 - a. Root directory – Verify that your source server has at least 4 GB of free disk space on the root directory (/)
 - b. RAM – Verify that your source server has at least 300 MB of free RAM to run the AWS Replication Agent.
 - 3. [Linux Installation Requirements](#)
 - 4. [Windows Installation Requirements](#)
- iii. [AWS Replication Agent installation instructions](#) – You must install the AWS Replication Agent on each source server that you want to add to Elastic Disaster Recovery. Agent installation is composed of the following steps:
 - 1. [Generating the required AWS credentials](#)
 - 2. [Installation instructions](#)
 - 3. [Installing the Agent on a secured network](#)
- d. [Configure Launch Settings](#). After you have added your source servers to the Elastic Disaster Recovery Console, you will need to configure the launch settings for each server. The launch settings are a set of instructions that determine how a recovery instance will be launched for each source server on AWS. You must configure the launch settings prior to launching test or recovery instances. You can use the default settings or configure the settings to fit your requirements.
 - i. Note: Change the "Copy source" option to "no" to ensure the launch does not fail.
 - ii. Note: You can change the launch settings after a drill or recovery instance has been launched. You will need to launch a new drill or recovery instance for the new settings to take effect.

- iii. Note: Windows source servers need to have at least 2 GB of free space to successfully launch a recovery instance.
- iv. Note: Take into consideration that once a drill instance is launched, actual resources will be used in your AWS account and you will be billed for these resources. You can terminate the operation of launched recovery instances once you verify that they are working properly.
- e. [Launch a drill instance](#). It is crucial to drill the recovery of your source servers to AWS prior to initiating a recovery in order to verify that your source servers function properly within the AWS environment. It is a best practice to have a VPC configured specifically for testing, that has no connection to your on-premises environment.
- f. [Launch a recovery instance](#)

Resources

- **Network diagrams:** <https://docs.aws.amazon.com/drs/latest/userguide/Network-diagrams.html>
- **Network setting preparation**
 - **Staging Area Subnet:** Before setting up Elastic Disaster Recovery you should create a subnet which will be used by Elastic Disaster Recovery as a staging area for data replicated from your source servers to AWS. You must specify this subnet in the replication settings template. [Learn more about replication settings](#).
 - **Network Requirements:** The replication servers launched by Elastic Disaster Recovery in your Staging Area Subnet need to be able to send data over TCP port 443 to the Elastic Disaster Recovery API endpoint at <https://drs.{region}.amazonaws.com/>. Replace “{region}” with the AWS Region code you are replicating to, for example “us-east-1”. The source servers on which the AWS Replication Agent is installed need be able to send data over TCP port 1500 to the replication servers in the Staging Area Subnet. They also need to be able to send data to Elastic Disaster Recovery's API endpoint at <https://drs.{region}.amazonaws.com/>. Replace “{region}” with the AWS Region code you are replicating to, for example “us-east-1”
 - **Operational Subnet:** Drill and recovery instances are launched in a subnet you specify in the Amazon EC2 Launch Template associated with each source server. The Amazon EC2 Launch Template is created automatically when you add a source server to Elastic Disaster Recovery.