



Migrate workloads to the AWS Cloud by using CloudEndure Migration

AWS Prescriptive Guidance



AWS Prescriptive Guidance: Migrate workloads to the AWS Cloud by using CloudEndure Migration

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Introduction	1
Architecture	3
Preparing your environment	6
Set up your CloudEndure account and users	6
Generate and use AWS Cloud credentials	8
Configure your source and target network	9
Network requirements	9
Communication over TCP port 443	11
Proxy servers	11
Migrating your workload	13
Define replication settings for AWS	13
Install CloudEndure Agent on source machines	14
Configure the target system Blueprint	15
Testing the migration	17
Cutting over to AWS	19
Best practices	21
FAQ	23
Next steps	26
Additional resources	27
Related guides and patterns	27
References	27
Tools	27
Document history	28

Migrate workloads to the AWS Cloud by using CloudEndure Migration

Chandra Sekhar Yaratha and Harnath Valeti, Amazon Web Services

Chandra Sekhar Yaratha and Harnath Valeti, Consultants, AWS Professional Services

November 2020 ([last update](#): January 2021)

This guide describes the process for migrating applications from a physical, virtual, or cloud-based infrastructure to the Amazon Web Services (AWS) Cloud by using CloudEndure Migration. It includes best practices and troubleshooting tips, which are also covered in the following video.

CloudEndure Migration is an agent-based tool that rehosts your applications on AWS. It supports self-service, highly automated, lift-and-shift migrations with minimal business disruption. You install the CloudEndure Agent on your source machines. The Agent replicates your applications and data in a staging area on AWS. After the initial replication, the CloudEndure Agent tracks and migrates changes from your source environment to the target staging area by using asynchronous, block-level data replication, without causing downtime or affecting performance.

You can use CloudEndure's self-service, web-based console to configure your target environment, to check for compability issues, and to validate that your applications are fully operational on AWS before you cut over. After you rehost your workloads on AWS, you can restructure, rewrite, and re-architect your applications and take advantage of AWS services.

CloudEndure Migration supports both Microsoft Windows and Linux-based workloads (see [supported versions](#)). For more information, see the [CloudEndure Migration website](#).

This guide is for project managers, system administrators, and migration architects who are planning to migrate their workloads to AWS. It describes the migration process, including preparing your source and target environments, configuring network settings, migrating your workloads, testing, and cutting over. If you're planning a large-scale migration that involves rehosting thousands of servers, you can further automate tasks by using CloudEndure Migration Factory (CEMF).


Note

[AWS Application Migration Service \(MGN\)](#) is the primary migration service recommended for lift-and-shift migrations to the AWS Cloud. Customers who currently use CloudEndure

Migration or AWS Server Migration Service (AWS SMS) are encouraged to switch to MGN for future migrations.

Architecture

CloudEndure Migration simplifies, expedites, and automates large-scale migrations to AWS. Continuous data replication takes place in the background, without application disruption or performance impact, which ensures that data is synchronized in real time and minimizes cutover windows. When you initiate migration cutover, CloudEndure runs a highly automated machine conversion and orchestration process, which reduces the potential for human error. After migration, even the most complex applications and databases run natively on AWS, without compatibility issues and with minimal IT skills necessary. The following diagram illustrates the migration process.



Benefits of using CloudEndure Migration include:


- **Easy migration** – You can run complex, large-scale migration projects rapidly, regardless of the application type, while significantly reducing risk.
- **Increased uptime** – You can maintain normal business operations throughout the replication process. CloudEndure Migration copies source machines continuously, without taking snapshots or writing any data to disks. This means that there is minimal performance impact and no need to reboot machines. Continuous replication also makes it easy to conduct non-disruptive tests and shortens cutover windows.

- **Reduced costs** – CloudEndure Migration is a single tool for migrating any application or database from any source infrastructure on supported operating systems to AWS. You can migrate legacy applications, third-party applications, and line-of-business applications. There is no need to invest in specialized cloud development, operating system or application-specific skills, or significant IT resources, which results in greatly reduced operational costs.

Migrating your workloads by using CloudEndure Migration involves four phases of activities:

1. **Preparing your environment.** Includes setting up your CloudEndure account, creating AWS credentials, and configuring your network.
2. **Migrating your workload.** Includes installing CloudEndure Agents and replicating your source environment in the AWS staging area.
3. **Testing the migration.** Includes verifying the target machine settings and validating that the target machines are operating correctly.
4. **Cutting over to AWS.** CloudEndure Migration automatically converts your machines to run natively on AWS.

These phases are illustrated in the following diagram and described in detail in the following sections.



Preparing your environment

Before you use CloudEndure Migration, you have to make sure that your source and target environments are prepared and configured correctly. These preparations include setting a CloudEndure account, creating AWS credentials, defining network requirements, and opening the correct ports.

Set up your CloudEndure account and users

Before you start your migration, follow these steps:

1. Subscribe to CloudEndure Migration in [AWS Marketplace](#), where it's available as a SaaS contract directly through your existing AWS account.
2. Create an account in CloudEndure. For instructions, see [Registering to CloudEndure Migration](#) in the CloudEndure documentation.
3. Add users. The account owner and account administrators can manage users and projects in the CloudEndure User Console. For information about inviting and managing users, see [User Management](#) in the CloudEndure documentation.
4. Create and configure a CloudEndure project. Log in to the [CloudEndure User Console](#) and follow the steps provided in [Working with Projects](#) in the CloudEndure documentation.

For project type, select **Migration**.

Create New Project

Project name
Migration Immersion Day

Project type
Migration

Target infrastructure
Amazon Web Services

License

➔ **MIGRATION**

ID: [REDACTED]

Expires: 4/20/2023

Usage: 9996 licenses remaining of 10000 total licenses

CANCEL **CREATE PROJECT**

Project type, Target infrastructure and License cannot be changed after Project is created

Note

Make sure that you have sufficient [migration licenses](#) for your project.

Generate and use AWS Cloud credentials

Before you install CloudEndure Agents, you must create the AWS Identity and Access Management (IAM) credentials required to configure the CloudEndure User Console, and collect the access key and secret access key for the CloudEndure users in your organization. Follow these steps:

1. Sign in to your AWS account.
2. Create a policy for CloudEndure in your AWS account, using the code from <https://docs.cloudendure.com/Content/IAMPolicy.json>. For more information about this step, see the [CloudEndure documentation](#).
3. Create a new IAM user and generate AWS credentials, which consist of an access key ID and a secret access key. For more information, see the [CloudEndure documentation](#) and the [IAM documentation](#).
4. In the CloudEndure User Console, choose **Setup & Info, AWS Credentials**.

CloudEndure
An AWS Company

Migration Immersion Day Migration | PROJECT NOT SET UP

Dashboard
Machines
Job Progress
Event Log
Setup & Info

LICENSES **AWS CREDENTIALS** REPLICATION SETTINGS OTHER SETTINGS

Staging Area Account Credentials

In order for the service to work, you need to provide credentials for Amazon Web Services.

- 1 The Project must have [these permissions](#)
- 2 Why we need these credentials
- 3 Find out where to get these credentials
- 4 In case you'd like to use CloudEndure together with the AWS Migration Hub and you have not selected your Migration Hub Home Region yet, you can select it [here](#).

AWS Access Key ID

AWS Secret Access Key

SAVE


5. Add the access key ID and secret access key you created in the previous step into the corresponding fields and save.

Configure your source and target network

In the target environment on AWS, prepare a subnet to use as the CloudEndure staging area. Use a private subnet if you want to isolate your application. You can use an existing subnet or create a new one. This subnet is used to launch the CloudEndure replication servers, which receive replicated data from the source machines.

In general, there are three points of contact between CloudEndure components and the external network:

- The CloudEndure Agent needs to communicate with the CloudEndure Service Manager.
- The CloudEndure Agent needs to communicate with the CloudEndure replication servers.
- CloudEndure replication servers need to communicate with the CloudEndure Service Manager and Amazon Simple Storage Service (Amazon S3).



Network requirements

Client	Server	Description
Configure your source and target network		

Agent	Outbound: TCP 1500	Replication server(s) (private/public network)	Inbound: TCP 1500	Production instance status and data (the actual data replication stream)
Agent	Outbound: TCP 1500	Management (public network)		• REST APIs used during Agent installation • Agent monitoring • Statistics for Agents
Replication server(s)	Outbound: TCP 443	Management (public network)		• Statistics for replication servers • Replication server logs • Replication server API

To prepare your network for running CloudEndure's solutions, you need to set the following connectivity settings:

- Communication over TCP port 443:
 - Between the source machines and CloudEndure Service Manager
 - Between the staging area and CloudEndure Service Manager
- Communication over TCP port 1500:
 - Between the source machines and the staging area

Note

For information about replication from a private subnet or server behind a NAT gateway or instance, see the [CloudEndure documentation](#).

Communication over TCP port 443

Add the following IP addresses and URLs to your firewall:

- CloudEndure Service Manager IP address (required for using CloudEndure software)
 - 52.72.172.158 (main service of console.cloudendure.com)
 - 52.53.92.136
- Amazon S3 service URLs (required for downloading CloudEndure software)
 - s3.amazonaws.com
 - s3.us-west-1.amazonaws.com
 - s3.eu-west-1.amazonaws.com

Proxy servers

If you want to use proxy servers in your environment, follow this guidance for CloudEndure to operate effectively:

- Make sure to set the proxy in replication settings, either as an IP address or as a fully qualified domain name (FQDN).
- Configure the proxy allow list for the CloudEndure Agent:
 - In Windows, the CloudEndure Agent runs as a system service, so make sure that the system account is part of the proxy allow list.
 - In Linux, the CloudEndure Agent creates a Linux user (named cloudendure) to run commands, so make sure that this user is part of the proxy allow list.
- Set the proxy:
 - In Windows, open the Command Prompt window as an administrator, and run the command:

```
setx https_proxy https://<proxy-ip>:<proxy-port>/ /m
```

- In Linux, change to **sudo** user and run the commands:

```
$ export http_proxy=http://server-ip:port/  
$ export http_proxy=http://127.0.0.1:3128/  
$ export http_proxy=http://proxy-server.mycorp.com:3128/  
(make sure to include the trailing slash "/")
```

Optionally, if you want to keep the environment variables for Windows, follow these steps:

1. In Windows Control Panel, choose **System and Security, System, Advanced system settings**.
2. In the **Advanced** tab of the **System Properties** dialog box, choose **Environment Variables**.
3. In the **System variables** section of the **Environment Variables** dialog box, choose **New** to add the https_proxy environment variable, or **Edit** if the variable already exists.
4. In the **Variable value** field of the https_proxy variable, enter https://
PROXY_ADDR:PROXY_PORT/ and then choose **OK**.
5. If the CloudEndure Agent was already installed, restart the service.

Migrating your workload

To migrate your workload to AWS, you use CloudEndure's web-based console to configure your source and target environments and to set up your replication servers. In this phase, you also install the CloudEndure Agent on your source machines, and specify a Blueprint for creating and configuring your target machine. With these components and configurations in place, replication from the source environment to the lightweight staging area can begin.

Define replication settings for AWS

To replicate data into or within the target environment, CloudEndure uses replication server(s) to provide a lightweight replication staging area on AWS. These servers are automatically launched into the subnet you created previously in the [Configure your source and target network](#) section.

The CloudEndure User Console address (console.cloudendure.com) is used for all cloud and on-premises platforms, except for specific cases.

In the CloudEndure User Console, you use the **Replication Settings** tab to define your source and target environments, and to configure the default replication servers in the staging area of the target infrastructure. For step-by-step instructions, see [Defining Replication Settings for AWS](#) in the CloudEndure documentation.

CloudEndure
An AWS Company

Migration Immersion Day Migration | PROJECT NOT FULLY SET UP PROJECT ACTIONS...

Dashboard
Machines
Job Progress
Event Log
Setup & Info

LICENSES AWS CREDENTIALS **REPLICATION SETTINGS** OTHER SETTINGS

Migration Source
Other Infrastructure

Migration Target
AWS - Amazon S3 - S3 Glacier Deep Archive - Glacier Vault

Replication Servers
Replication Servers are small machines used to facilitate data replication. They are automatically launched in the [Staging Area](#) and discarded when no longer needed. [Learn about failing back to "Other Infrastructure"](#).

Choose the Replication Server instance type:
Default

Choose the Converter instance type:
Default (m5.xlarge/m4.xlarge, based on AZ)

Choose if each source machine should have a dedicated Replication Server:
☐ Use dedicated Replication Servers [learn more](#)

Choose the default disk type to be used by the Replication Servers (SSD disks are faster; HDD disks cost less; [learn more](#))
Use fast SSD data disks

Choose the subnet where the Replication Servers will be launched:
subnet-... (TargetPublic) from vpc-... (Target)

Choose the Security Groups to apply to the Replication Servers:
Default CloudEndure Security Group

Choose how data is sent from the CloudEndure Agent to the Replication Servers:
☐ Use VPN or DirectConnect (using a private IP) [learn more](#)

Define whether to route communication from the Replication Server via a proxy:


Install CloudEndure Agent on source machines

After you've prepared your environment and configured your replication settings, you can install CloudEndure Agent on the machines in your source environment. The Agent replicates your applications and data in the staging area on AWS. Follow these steps:

1. In the CloudEndure User Console, choose **Machines, Machine Actions, Add Machines**.
2. Locate the Installation Token that was provided for you when your CloudEndure account was activated. (To locate your tokens, choose **Help, How to Add Machines** in the CloudEndure User Console.)
3. Download the Agent installer file to each source machine, and then use the command provided in the **How to Add Machines** pane to run the installer on each source machine.

For more information, see [Installing the CloudEndure Agents](#) in the CloudEndure documentation.

When the Agent is installed, the source machine appears in the CloudEndure User Console, in the **Machines** tab, and the machine undergoes initial data replication (Initial Sync).



Use the **Machines** dashboard to monitor the Initial Sync process, which includes these steps:

1. Creating firewall rules.
2. Creating the replication server.
3. Booting the replication server.
4. Resolving the CloudEndure Service Manager address in the replication server.
5. Authenticating the replication server with the CloudEndure Service Manager.
6. Downloading the CloudEndure replication software to the replication server.
7. Creating staging disks.
8. Attaching the staging disks to the replication server.
9. Pairing the CloudEndure Agent with the replication server.
10. Establishing communication between the CloudEndure Agent and the replication server.

When the Initial Sync is complete, the **Machines** dashboard shows the final status as **Continuous Data Replication**. This indicates that all changes in the source server are being synchronized with the CloudEndure replication server.

For more information, see [Monitoring Your Solution](#) in the CloudEndure documentation.

Configure the target system Blueprint

A Blueprint provides a set of instructions for creating and launching a target machine for a specified source machine. To configure a source machine's Blueprint, choose the machine name from the **Machines** dashboard to access the **Machine Details** pane, and choose the **Blueprint** tab from the menu on the right.

Enter the following values in the **Blueprint** tab and then choose **Save Blueprint**.

Key	Value
Machine Type	t3.medium
Launch Type	On Demand
Subnet	<i>Target subnet</i>
Security Groups	<i>Target security group</i>
Private IP	Create New
Public IP (ephemeral):	No
Disks	SSD


For more information about these settings, see [Configuring the Target Machine Blueprint](#) in the CloudEndure documentation.

Testing the migration

Before you migrate your source machines into the target infrastructure, you should test your CloudEndure Migration configuration. The **Test Mode** workflow launches and runs a target machine in the target infrastructure for the source machine you selected for testing. By testing your migration configuration, you can verify that your source machines are working properly in the target environment. The CloudEndure User Console displays the test results. You can run **Test Mode** after the Initial Sync stage has been completed.

To test your migration, follow these steps:

1. Confirm that the source machine you want to test is in Continuous Data Replication (CDR) mode or that its status is **Ready for testing**.
2. Configure the target machine with Blueprint. Verify that the subnet of the target environment is isolated. This isolation is designed to prevent conflicts with the source environment.
3. Test the target machines:
 - a. On the **Machines** page, check the box to the left of each source machine you want to test.
 - b. **Choose Launch x Target Machine**, and then choose **Test Mode**.



1. When you receive the confirmation prompt, choose **Continue** to launch the target machines. You can monitor the launch process on the **Job Progress** tab.
2. Verify that the test completed successfully.


3. Test the target machines by choosing each machine's name, navigating to the **Target** tab, copying the public IP, and navigating to that IP.
4. Verify service configuration and other settings.
5. Verify that your network works as expected.
6. Validate that you can connect to your target machines by using Secure Shell (SSH) for Linux or Remote Desktop Protocol (RDP) for Windows, and perform acceptance tests for your application.

For more information about the testing process, see [Testing the Migration Solution](#) in the CloudEndure documentation.

Cutting over to AWS

When you've completed your test and you're ready to cut over to the target environment on AWS, follow these steps:

1. Confirm that the **Data Replication Progress** for the source machine is **Continuous Data Replication** mode, and that the **Live Migration Lifecycle** column displays the status **Tested**.
2. Verify the Blueprint configuration.
3. Schedule and perform the cutover:
 - a. On the **Machines** page, check the box to the left of each source machine you want to migrate.
 - b. **Choose Launch x Target Machine**, and then choose **Cutover Mode**.



- c. When you receive the confirmation prompt, choose **Continue** to launch the target machines. You can monitor the launch process on the **Job Progress** tab.
- d. Verify that the cutover completed successfully.

Job Progress

Type: cutover launch | Status: Completed Successfully | Started: 4/21/2020, 4:37:16 PM | Finished: 4/21/2020, 4:41:11 PM

4/21/2020, 4:39:20 PM Finished creating a replica for security group

4/21/2020, 4:39:20 PM Finished creating a replica for security group

4/21/2020, 4:39:20 PM Finished creating a replica for security group

4/21/2020, 4:39:21 PM Started creating a replica for instance ofbiz-web.onpremsim.env

4/21/2020, 4:39:36 PM Started creating a replica for instance wordpress-db.onpremsim.env

4/21/2020, 4:39:52 PM Started creating a replica for instance wordpress-web.onpremsim.env

4/21/2020, 4:40:14 PM Finished creating a replica for instance ofbiz-web.onpremsim.env

4/21/2020, 4:40:20 PM Finished creating a replica for instance wordpress-db.onpremsim.env

4/21/2020, 4:40:24 PM Started creating a replica for instance ofbiz-db.onpremsim.env

4/21/2020, 4:40:37 PM Finished creating a replica for instance wordpress-web.onpremsim.env

4/21/2020, 4:41:09 PM Finished creating a replica for instance ofbiz-db.onpremsim.env

4/21/2020, 4:41:11 PM **Job finished**



CLOSE

- e. Verify service configuration and other settings.
 - f. Verify that your network works as expected.
 - g. Validate that you can connect to your target machines by using SSH for Linux or RDP for Windows, and perform acceptance tests for your application.
 - h. Shut down your source machines.
4. After cutover is validated, uninstall the CloudEndure Agent by removing machines from the CloudEndure User Console.
 - a. On the **Machines** page, check the box to the left of each source machine you want to remove.
 - b. From the **Machine Actions** menu, choose **Remove x Machines from This Console**. It takes up to 60 minutes for CloudEndure Migration to clean up the replication instances and volumes in the staging area.
 - c. When all Agents have been uninstalled, delete the virtual private cloud (VPC) for the staging area. This deletes all the AWS resources that you created for replication.

For more information about the cutover process, see [Performing a Migration Cutover](#) in the CloudEndure documentation.

Best practices

The following video covers best practices and troubleshooting tips for using CloudEndure Migration.

To ensure project success:

- Train a field technical team, and assign a CloudEndure subject matter expert (SME).
- Share project timelines with CloudEndure.
- Monitor data replication progress and report any issues in advance.
- Perform a test for every machine in advance, and report issues to CloudEndure.
- Coordinate cutover windows with CloudEndure in advance.

Follow these best practices to implement a smooth and efficient migration process:

- Create a service account in the CloudEndure User Console instead of using individual administrator accounts.
- Make sure that you have sufficient migration licenses for your project.
- Install CloudEndure Agent on your source machines at least three weeks before your planned cutover. This will give you enough time to test and help identify any issues in advance.
- Install CloudEndure Agents on the source machines by using the `--no-replication` argument, to prevent the replication from starting automatically.
- When scheduling your cutover, allow enough time for data replication to complete and for all necessary testing to be carried out.
- Complete cutover within 90 days after you install the CloudEndure Agent. Migration licenses expire after 90 days.
- Do not perform any reboots on the source machines before a cutover.
- Perform a test cutover at least one week before you plan to migrate your source machines. This will help you identify potential problems and solve them before the actual cutover takes place. After performing the test launch:
 - Validate connectivity to your target machines (using SSH for Linux or RDP for Windows), and perform acceptance tests for your application.
 - Verify that the subnet of the target site is isolated, to prevent conflicts with the source site.
 - Verify that your network works as expected.

- Verify service configuration and other settings.
- For **Network Bandwidth Throttling**, select **Disabled**. Enable this option only if you want to control the amount of bandwidth used for replication traffic.
- Remove the source machines from the CloudEndure User Console after cutover is complete, to clean up the staging area, remove the replication resources that are no longer needed, and reduce your costs.
- If you're replicating a large instance, use a dedicated replication server to accelerate replication.
- After final cutover, verify that the launched Amazon Elastic Compute Cloud (Amazon EC2) instance passes health checks (*2/2 checks passed* status) on AWS.

FAQ

This section provides answers to commonly raised questions about using CloudEndure Migration to migrate your workloads. For additional information, see the [FAQ](#) on the CloudEndure website.

Q. Can I use my existing Windows licensing with CloudEndure Migration?

A. You have a variety of options for using new and existing Microsoft software licenses on the AWS Cloud.

- You can purchase Amazon EC2 or Amazon Relational Database Service (Amazon RDS) license-included instances to get new, fully compliant Windows Server and SQL Server licenses from AWS.
- You can use your existing licenses on AWS with Amazon EC2 Dedicated Hosts, Amazon EC2 Dedicated Instances, or EC2 instances with default tenancy by using [Microsoft License Mobility through Software Assurance](#). Whether you want to purchase new Microsoft licenses, or utilize existing ones, you can rely on AWS to run your Microsoft software.

CloudEndure converts Microsoft Windows licenses to AWS Windows licenses and activates them by using Microsoft Key Management Service. If license activation fails, follow the [instructions on the AWS support website](#) to resolve the issue.

Q. Can I use my existing Linux licenses with CloudEndure Migration?

A. When you migrate your Linux machines to AWS with CloudEndure Migration, you must supply your own Linux OS licenses (BYOL). That is, AWS will not assign its own Linux OS licenses automatically.

Q. What is the required bandwidth for TCP port 1500 to replicate data?

A. The required bandwidth for transferring the replicated data over TCP port 1500 should be based on the write speed of the participating source machines. The minimum bandwidth is the sum of the average write speed of all replicated source machines.

For example, if you are replicating two source machines that have write speeds of 5 megabytes per second (MBps) and 7 MBps, the recommended bandwidth should be at least 12 MBps.

Q. Can I change the default disk type to SSD instead of Provisioned SSD in the Blueprint for the target instance?

A. You can change the default disk type, but this affects the first boot time of the target machine. The target machine will boot much more slowly the first time on AWS. Sometimes, this process is so slow that some Windows processes will time out and network drivers won't be activated during bootup time.

Q. How can I route replication traffic when I have multiple leased lines in the source? Is there configuration setting I can use in CloudEndure? **A.** CloudEndure uses standard routing. If you configure the router to push the data to the IP addresses to the staging area over the second link, this is what it'll do.

Q. What compression mechanism does CloudEndure use during replication? **A.** CloudEndure uses [Lempel-Ziv-Welch \(LZW\) compression](#) during transit, which results in 60-70% compression, depending on the type of data.

Q. What does the CloudEndure machine conversion server do?

A. The machine conversion server runs in the target infrastructure and converts the disks to boot. Specifically, this server makes bootloader changes, injects hypervisor drivers, and installs cloud tools. The CloudEndure machine conversion process typically takes less than a minute.

Q. What are the differences between conversion servers and replication servers?

A. Here's a list of similarities and differences between the two server types:

- Replication servers run on Linux, and conversion servers (for Windows machines) run on Windows.
- When CloudEndure launches replication servers, it automatically places Windows conversion servers in the same subnet as the replication servers.
- Both conversion and replication servers have public IPs.
- Conversion servers use the same security groups as replication servers.
- The conversion server must be able to access the CloudEndure Service Manager.
- The conversion server machines, just like replication servers, are managed automatically by CloudEndure. Any attempt to disrupt their automated functionality will result in failed conversions.

Q. Which AWS Regions does CloudEndure use for production and disaster recovery?

A. CloudEndure uses US East (N. Virginia) for production and US West (N. California) for disaster recovery. Data is copied from the production Region to the disaster recovery Region over a private AWS network.

Q. What is the effect of having a dedicated replication server?

A. The dedicated replication instance makes full use of available bandwidth to accelerate replication.

Next steps

Because CloudEndure Migration minimizes cutover downtime and provides continuous replication, it handles migration scenarios such as migrating your servers to AWS, deploying a block-level replication solution, or supporting real-time replication requirements much better than other migration tools.

If your enterprise is committed to a cloud migration journey and you wish to achieve a range of business benefits by migrating existing workloads to AWS, learn more about the [AWS Migration Acceleration Program \(MAP\)](#), which provides consulting support, training, and services credits.

Before you migrate your production workloads, we recommend that you run a proof of concept (POC) to understand the CloudEndure Migration tool, infrastructure requirements, data transfer speeds, and other migration considerations.

Additional resources

Related guides and patterns

- [CloudEndure Migration patterns](#)
- [Automating large-scale server migrations with CloudEndure Migration Factory](#)
- [Migrating Oracle databases to the AWS Cloud](#)
- [Migrating SQL Server databases to the AWS Cloud](#)
- [AWS Prescriptive Guidance website](#)

References

- [Migrating workloads across AWS Regions with CloudEndure Migration](#)
- [Migration Best Practices](#) (CloudEndure documentation)
- [Troubleshooting Playbook](#) (CloudEndure documentation)
- [CloudEndure documentation home page](#)

Tools

- [CloudEndure API documentation](#)
- [CloudEndure API](#)
- [AWS CloudEndure Migration Factory solution](#)
- [AWS Migration Hub](#)
- [AWS Migration Evaluator](#)
- [CloudEndure Migration Training - Technical](#)
- [CloudEndure Disaster Recovery Technical Training](#)

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an [RSS feed](#).

Change	Description	Date
Added video	We added a video that covers best practices and troubleshooting tips.	January 27, 2021
Initial publication	—	November 23, 2020