

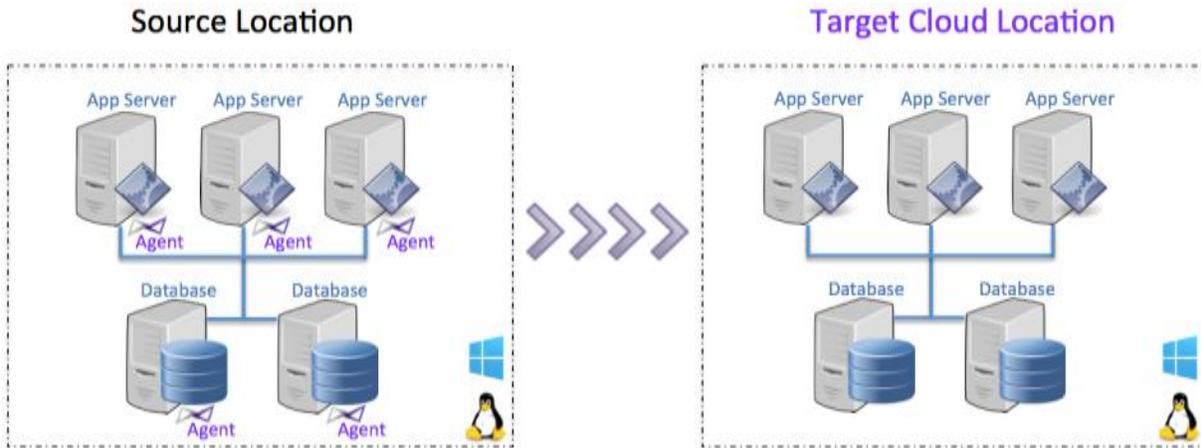


Cloud**Endure**

# How CloudEndure Works

Whitepaper

CloudEndure provides software based Cloud Workload Mobility solutions that are specifically designed to migrate or copy workloads to and within cloud environments, including use cases such as Live-Migration to the Cloud and Disaster Recovery to and in the Cloud. When required, you may migrate or copy your live workloads including their up-to-date data to a target cloud location.



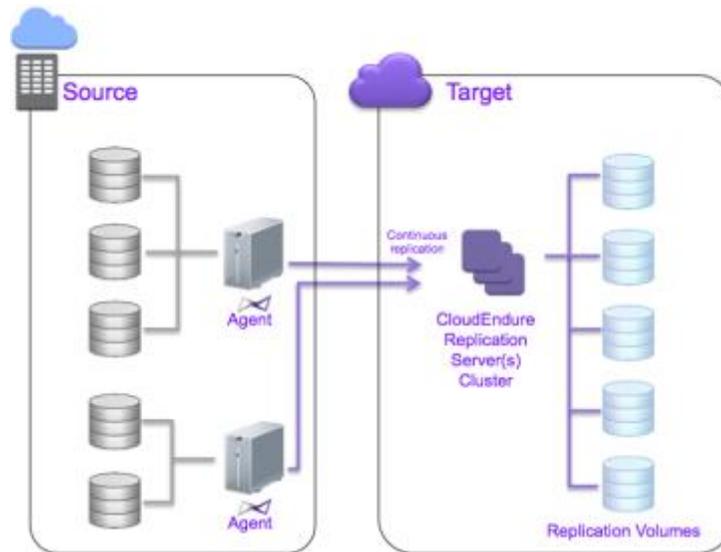
Real-time replication  
**From ANY** infrastructure  
 To any supported Private/Public Cloud



## How the Technology Works

CloudEndure replicates workloads to a target cloud location, where a fully functioning, up-to-date copy of the application can be spun up on demand. A combination of continuous block level replication of all volumes attached to all machines, as well as constant replication of the entire workload application stack (machine types, number and sizes of attached volumes, network topology, and so on), results in fail-safe mobility of your workloads to and within the cloud. The target workload is always up-to-date, as it is not snapshot based, but the product of continuous streaming of real-time data (CDP).

## Volume Replication



CloudEndure provides a low-footprint agent that the customer installs on all of the machines that are a part of the workload. The agent performs two tasks:

- Initial block-level read of the content of any volume attached to the machine and one-time replication of the content to the target cloud location (chosen by the customer).  
This one-time activity may take between several minutes and several days, depending on the amount of data to be copied and the available bandwidth between the source location and the target cloud location.
- Real-Time monitoring of all block-level modifications to any volumes attached to the machine.  
Once initial replication is complete, CloudEndure's replication servers, located in the target cloud location, are continuously synchronized to the latest block-level modifications, providing near-zero RPO (Recovery Point Objective). Each time any volume content is modified by an application running on the machine, the agent sends the change, in real-time, to CloudEndure's replication server(s) located in the target cloud location (located in either the customer's or CloudEndure's cloud account). The replication server immediately writes the change to the replica volume.  
The monitoring activity starts in parallel to the initial copying, so any changes to the content of the volumes during the initial copy are also replicated, using CloudEndure's proprietary algorithm to maintain consistent copies of the volumes. The data sent during this task is compressed and encrypted using AES-256 bit.

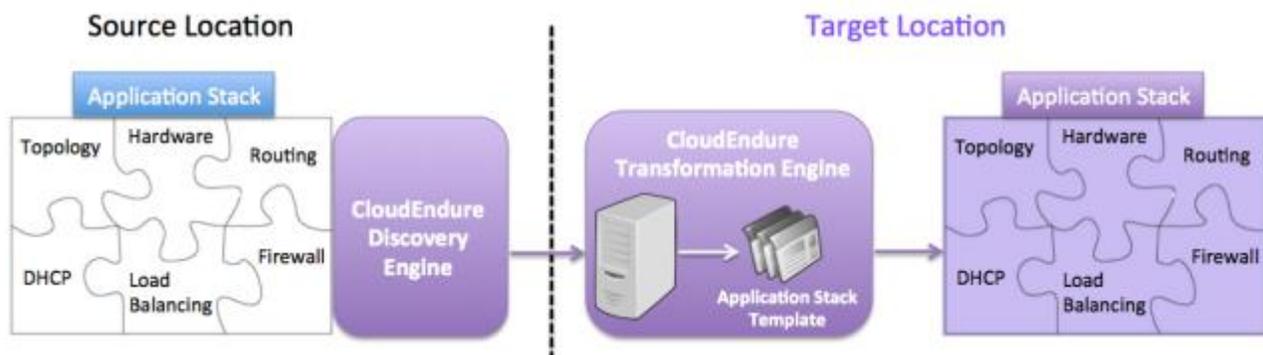
CloudEndure servers maintain consistency across all volumes attached to an instance, including multiple volumes that are a part of a RAID array.

CloudEndure's replication mechanism is not snapshot-based, and uses CDP, meaning that the target cloud location always remains in sync with the source and no downtime is required when migrating or copying the workload. Using continuous sync also eliminates potential performance issues caused by snapshot-based replication. To read more on this subject, please refer to the "CloudEndure - CDP vs Alternatives" whitepaper.

## 1-Minute Image Conversion

When replicating machines across regions within a certain cloud, the machines can boot natively in the target region, as there are no significant changes in infrastructure. Things are not as simple however, when replicating machines between different infrastructures, whether from on-premise (physical/virtual) to the cloud, or between different cloud infrastructure providers. In such cases, the replicated machines cannot boot and work out of the box, due to infrastructure differences in hypervisors, unique cloud tools and other variations. CloudEndure addresses this by using its proprietary machine conversion technology, which handles hypervisor and various OS configuration changes, boot process changes, OS activation, injection of cloud tools etc. The automated machine conversion process takes ~30 seconds and ensures that any Windows/Linux server, coming from any source physical/virtual/cloud infrastructure will natively boot and run transparently in the target cloud of your choice.

## Application Stack Orchestration



The application stack is the component of the workload that connects its multiple machines into a working application. It includes the network topology, firewall configuration, load balancing settings, hardware components and more. This application stack is often very complex and in order to successfully migrate or copy a workload to a different location, the entire application stack also needs to be replicated correctly, without any gaps. Application stacks naturally look and behave differently between on-premise and cloud environments, and between different cloud environment infrastructures. Furthermore, application stacks differ even between different regions of the same cloud vendor. CloudEndure automatically replicates application stacks in two steps:

### Automatic Application Stack Discovery

CloudEndure uses APIs of source cloud-based and virtualized environments that it supports to automatically discover the entire application stack of the source workload. CloudEndure periodically repeats this process to make sure that any changes made to the application stack are propagated by CloudEndure to the target cloud location and will be taken into account when CloudEndure re-creates the workload in that location. For bare-metal source environments, CloudEndure lets the user to define the target application stack manually instead.

## Automatic Application Stack Transformation

To make sure that the workload is created properly in the target cloud location, its application stack converted by CloudEndure to the format of that location, taking into account the location's specific properties and adjusting it for the dynamic conditions of the location at that moment (for example, making sure that that location has the capacity to create the necessary resource and if it doesn't, modifying the resources accordingly). CloudEndure can transform an application stack to any of the target cloud environments it supports.

## Smart Resource Allocation

The uniqueness of CloudEndure is that it spins-up instances and provisions services on the target cloud location only when they are needed. During the replication process, all the data (including the content of the volumes and the application stack) is stored continuously. However, the replicas of the instances that comprise your application are not spun-up yet and no replica services are provisioned.

Only when the user explicitly requests the creation of the replica workload, do the following occur:

- The replica instances are spun-up and started in the target location.
- The replica services are provisioned to connect the instances into a functioning workload.

## Platform Support

CloudEndure allows replication of any machine, whether bare-metal, virtualized or cloud-based.

Supported operating systems:

- Ubuntu v10.04 or higher, 32/64 bit
- Redhat Enterprise Linux v5 or higher, 32/64 bit
- CentOS v5 or higher, 32/64 bit
- SUSE Linux v11 or higher
- Debian Linux 8 or higher
- Oracle Linux v5 or higher
- Amazon Linux 2013.03 or higher, 32/64 bit
- Microsoft Windows Server 2003, 2003 R2, 2008, 2008 R2, 2012, 2012 R2, 2016, 32/64 bit

## High Availability and Security

CloudEndure was designed and built by high availability, network and security experts with dozens of years of experience. The service is infinitely scalable, always on, and assures secure storage and transfer of all data.

For further questions, please contact [sales@cloudendure.com](mailto:sales@cloudendure.com)