

How CloudEndure Works

THE TECHNOLOGY BEHIND CLOUDENDURE'S DISASTER RECOVERY AND LIVE MIGRATION SOLUTIONS

CloudEndure offers Disaster Recovery and Live Migration Software-as-a-Service (SaaS) solutions. Both solutions utilize CloudEndure's mobility technology, which continually replicates workloads from any source infrastructure to a low-cost "staging area" (detailed below) in any target infrastructure, where an up-to-date copy of the workloads can be spun up on demand and be fully functioning in minutes.

In addition to a self-service, web-based Console with centralized management for all Disaster Recovery and Live Migration projects, CloudEndure provides APIs that enable developers to implement large-scale automation and other advanced capabilities.

CloudEndure Disaster Recovery

CloudEndure Disaster Recovery enables organizations to quickly and easily shift their disaster recovery strategy to use public clouds, private clouds, or existing VMware-based data centers, providing ultimate Hybrid IT Resilience flexibility. CloudEndure's disaster recovery solution utilizes block-level, Continuous Data Replication, which ensures that target machines are spun up in their most up-to-date state during a disaster or drill. Organizations can thereby consistently achieve sub-second Recovery Point Objectives (RPOs).

The Continuous Data Replication takes place in a low-cost "staging area," which reduces compute and storage footprint to a minimum. In the event of a disaster, CloudEndure triggers an automated system conversion process (p2c/v2c/c2c) and a scalable orchestration engine that can spin up thousands

of machines in the target infrastructure within minutes. This enables organizations to achieve Recovery Time Objectives (RTOs) of minutes. Accordingly, CloudEndure's disaster recovery solution provides the resilience of a warm standby solution at the low cost of a cold standby solution.

CloudEndure Disaster Recovery supports recovery from all physical, virtual, and hybrid cloud infrastructures into AWS, Azure, Google Cloud Platform (GCP), VMware, and OpenStack-based infrastructures.

Benefits of CloudEndure Disaster Recovery include:

- ➔ Average savings of 80% on total cost of ownership (TCO) compared to traditional disaster recovery solutions
- ➔ Sub-second Recovery Point Objectives (RPOs)
- ➔ Recovery Time Objectives (RTOs) of minutes
- ➔ Support of all application types, including databases and other write-intensive workloads
- ➔ Automated failover to target site during a disaster
- ➔ Point-in-time recovery, enabling failover to earlier versions of replicated servers
- ➔ One-click failback, restoring operations to source servers automatically
- ➔ Unlimited, non-disruptive disaster recovery drills

CloudEndure Live Migration

CloudEndure Live Migration provides migration to, across, and within any target infrastructure from any source infrastructure. The low-cost "staging area" enables pre-migration replication to occur without performance disruption or data loss. Automated

machine conversion and application stack orchestration ensure minimal cutover windows.

CloudEndure Live Migration supports migration to all on-premise environments, including VMware and all public and private clouds. The clouds we support include AWS (including GovCloud), Azure, Google Cloud Platform (GCP), OpenStack, Oracle Cloud, CloudStack, IBM Cloud, and more.

Benefits of CloudEndure Live Migration include:

- ➔ Cutover windows of minutes and no data loss
- ➔ 100% data integrity for all applications (including databases and legacy applications)
- ➔ Large-scale migrations with no performance impact
- ➔ Support for all source and target infrastructures
- ➔ Automated migration to minimize IT resources and project length

Continuous Data Replication

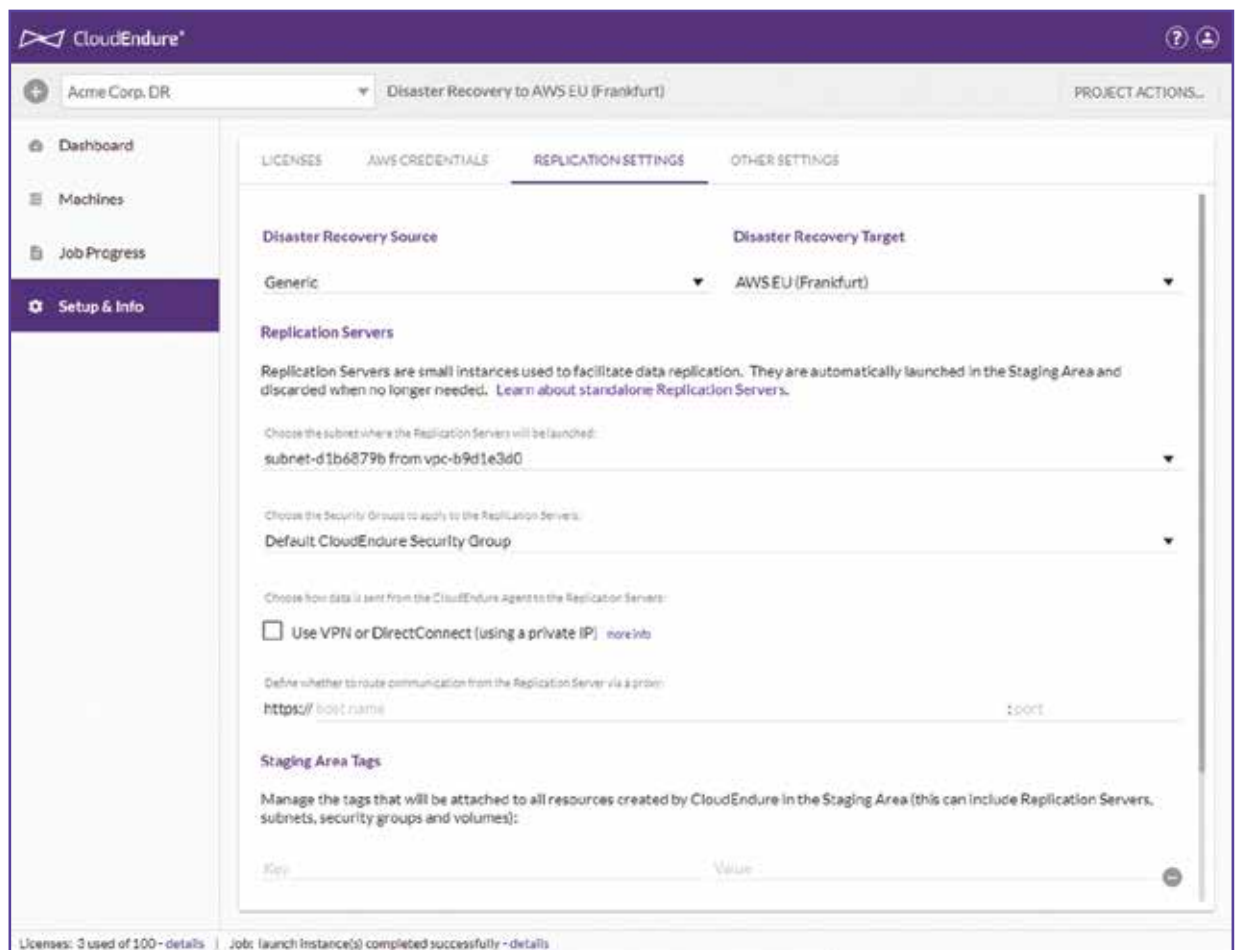
At the core of CloudEndure’s technology is a proprietary Continuous Data Replication engine, which provides real-time, asynchronous, block-level replication for Live Migration and Disaster Recovery.

CloudEndure replication is done at the OS level (rather than hypervisor or SAN level), enabling support of any type of source infrastructure:

- ➔ Physical machines, including both on-premises and co-location data centers
- ➔ Virtual machines, including VMware, Microsoft Hyper-V, and others
- ➔ Cloud-based machines, including AWS, Microsoft Azure, Google Cloud Platform, Oracle Cloud, OpenStack, and others

CloudEndure Console Replication Settings

Customers select disaster recovery source and target, subnets, VLANs, security groups, and more from the CloudEndure Console before beginning initial replication.



Once installed and activated, the CloudEndure agent begins initial replication, reading all of the data on the machines at the block level and replicating it to a low-cost “staging area” in the customer’s individual account in their preferred target infrastructure. Customers select their preferred target infrastructure as well as other replication settings such as subnets, VLANs, security groups, replication tags, and more, through the self-service, web-based CloudEndure Console.

The initial replication can take anywhere from several minutes to several days, depending on the amount of data to be replicated and the bandwidth available between the source and target infrastructures. No reboot is required nor is there any system disruption throughout the initial replication.

After the initial replication is complete, the source machines are continuously monitored to ensure constant synchronization, up to the last second. Any changes to source machines are

asynchronously replicated in real-time into the “staging area” in the target infrastructure.

Unlike snapshot-based replication, Continuous Data Replication enables customers to continue normal IT operations during the entire replication process without performance disruption or data loss.

For the disaster recovery use case, Continuous Data Replication allows customers to achieve sub-second Recovery Point Objectives (RPOs) as the data is always up-to-date and ready to be spun up as soon as a disaster strikes.

CloudEndure’s replication engine is application-agnostic, supporting all application types, including databases and legacy applications, all hardware configurations, and a wide array of operating systems and infrastructures. A partial list can be found below.

Applications*	Hardware Configurations*	Operating Systems*	Infrastructures*
Apache	Linux LVM	Amazon Linux	AWS
Apache Cassandra	Local Disks	CentOS	Azure
Apache Hbase	Oracle ASM	Debian	CloudStack
CouchDB	RAID Array	Kali	Google Cloud Platform
Microsoft Active Directory	SAN-Based Disks	Oracle Linux	IBM SoftLayer/Bluemix
Microsoft Dynamics CRM	Striped Disks	RedHat Linux	KVM
Microsoft Exchange		SUSE	Microsoft Hyper-V
Microsoft IIS		Ubuntu	OpenStack
Microsoft SharePoint		Windows Server 2003, 2008, 2012, 2016	Oracle Cloud
Microsoft SQL Server			Physical Servers
MongoDB			VMware
MySQL			Xen
NGINX			
openSUSE			
Oracle Database			
Oracle ERP			
Oracle Peoplesoft			
Pivotal CRM			
PostgreSQL			
Redis			
SAP CRM			
SAP ERP			
SAP S/4 Hana			

◀ **Applications, Hardware Configurations, Operating Systems, & Infrastructures Supported by CloudEndure (partial list)**

**As this is only a partial list, please [contact CloudEndure](#) if you do not see your application, hardware configuration, OS, or infrastructure in this table.*

Low-Cost “Staging Area” in Target Cloud

CloudEndure maintains ongoing replication of source machines into a low-cost “staging area” in the customer’s preferred target infrastructure. The “staging area” contains cost-effective resources automatically created and managed by CloudEndure to receive the replicated data without incurring any significant costs. These resources include a small number of VMs (each supporting multiple source machines), disks (one target disk for each replicating source disk), and snapshots.

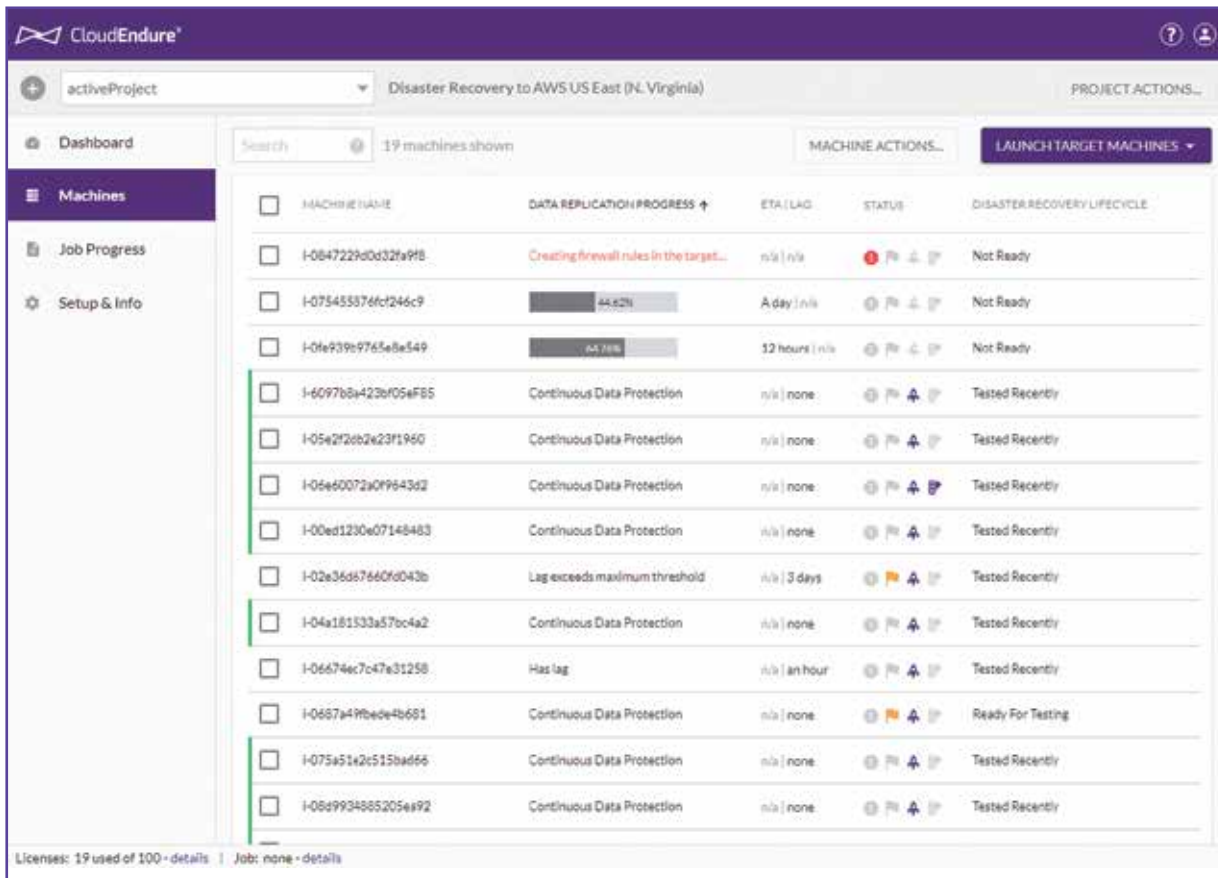
Whereas traditional disaster recovery solutions require duplicate hardware, compute, storage, networking, and software licenses in order to ensure data integrity and near-zero RPO, CloudEndure’s “staging area” eliminates the need for duplicate provisioning of resources and reduces the TCO for disaster recovery by an average of 80%. The more expensive recovery environment, which uses high-performance storage and actual compute to run applications, is only utilized when launched during a disaster or drill.

Automated Orchestration of Target Environment

CloudEndure’s orchestration engine automatically launches completely functional workloads in the target environment of the customer’s choice. This automated process includes cloning disks from the “staging area” to target networks and provisioning all the additional resources such as VMs, network interfaces, firewalls, etc.

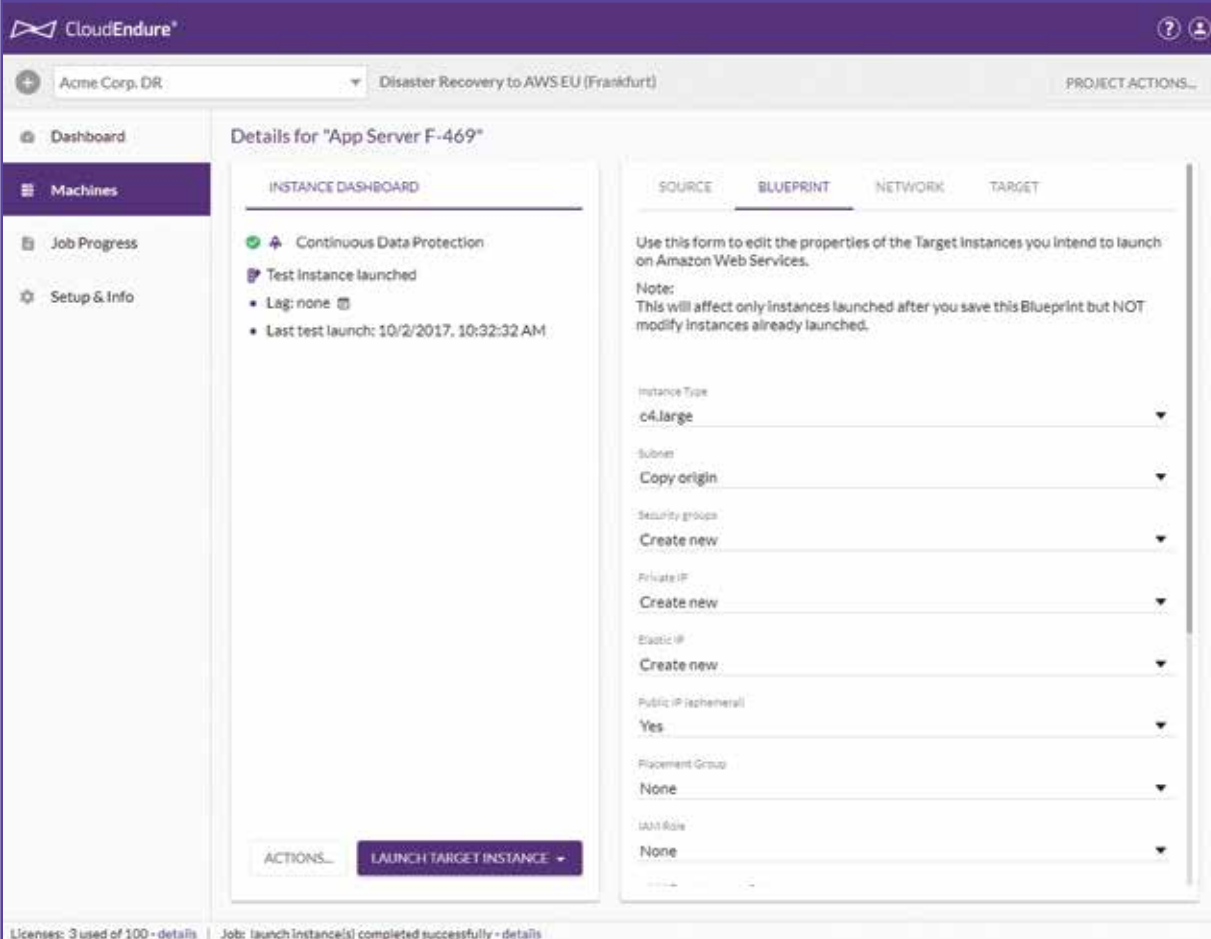
The entire orchestration process typically takes minutes. (Exact times vary depending on the OS and target environment.) CloudEndure’s orchestration process is not impacted by the number or size of volumes or the number of VMs.

For the disaster recovery use case, automated orchestration combined with machine conversion (detailed below) enable customers to achieve Recovery Time Objectives (RTOs) of minutes.



◀ **CloudEndure Console**
List of Machines & Data Replication Progress

The CloudEndure Console provides real-time information on the data replication status of each machine. Customers can add machines, conduct drills, and launch failovers and failbacks right from the Console.



CloudEndure

Acme Corp. DR

Disaster Recovery to AWS EU (Frankfurt)

PROJECT ACTIONS...

Dashboard

Machines

Job Progress

Setup & Info

Details for "App Server F-469"

INSTANCE DASHBOARD

- Continuous Data Protection
- Test Instance launched
 - Lag: none
 - Last test launch: 10/2/2017, 10:32:32 AM

ACTIONS... LAUNCH TARGET INSTANCE

BLUEPRINT

Use this form to edit the properties of the Target Instances you intend to launch on Amazon Web Services.

Note: This will affect only instances launched after you save this Blueprint but NOT modify instances already launched.

Instance Type: c4.large

Subnet: Copy origin

Security groups: Create new

Private IP: Create new

Elastic IP: Create new

Public IP (optional): Yes

Placement Group: None

IAM Role: None

Licenses: 3 used of 100 - details | Job: launchInstance(s) completed successfully - details

CloudEndure Console Blueprint Configuration Settings

Customers are able to select configuration settings for how machines will be provisioned (orchestrated) in the target infrastructure, including instance/machine type, subnet, VLAN(s), security groups, elastic IP, and more.

Customers are able to select configuration settings for how machines will be provisioned (orchestrated) in the target infrastructure, including instance/machine type, subnet, VLAN(s), security groups, elastic IP, and more, in the CloudEndure Console.

Automated Machine Conversion

When replicating machines across similar infrastructures, the replicated machines can boot natively in the target environment, as there are no significant differences in infrastructure. However, when replicating machines across dissimilar infrastructures, most disaster recovery solutions fail in making the replicated machines usable in the target infrastructure, due to infrastructure differences in hypervisors, drivers, and other variations. Without proper conversion, such transitions between physical machines, hypervisor variations or different clouds result in non-bootable target machines.

CloudEndure addresses this by using its proprietary machine conversion technology, which handles all hypervisor and OS configuration changes, boot process changes, OS activation, and installation of target infrastructure guest agents. The automated machine conversion process takes approximately 30 seconds and ensures that any Windows/Linux machine replicated from any source (physical/virtual/cloud) will natively boot and run transparently in the customer's preferred target location.

Automated Failback

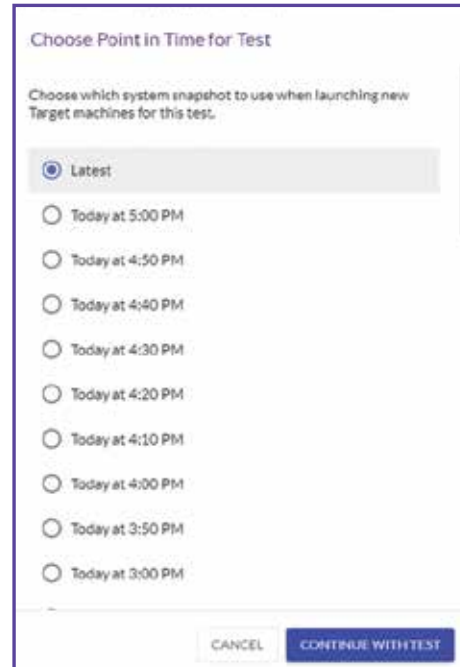
Once a disaster is over, CloudEndure provides automated failback to the source infrastructure. Because CloudEndure's failback technology also utilizes Continuous Data Replication, failback to source machines is rapid and no data is lost during the process. CloudEndure's automated failback supports both incremental and bare-metal restores.

Point-in-Time Recovery

In many disaster recovery use cases, the goal is to spin up the most up-to-date state of the source machines and continue operations as before. However, in cases of database corruptions, accidental system changes, ransomware, and other malicious attacks, customers may need to recover to previous consistent points in time. CloudEndure Disaster Recovery offers point-in-time recovery, which protects and recovers data and IT environments that have been corrupted.

Enterprise-Grade Security

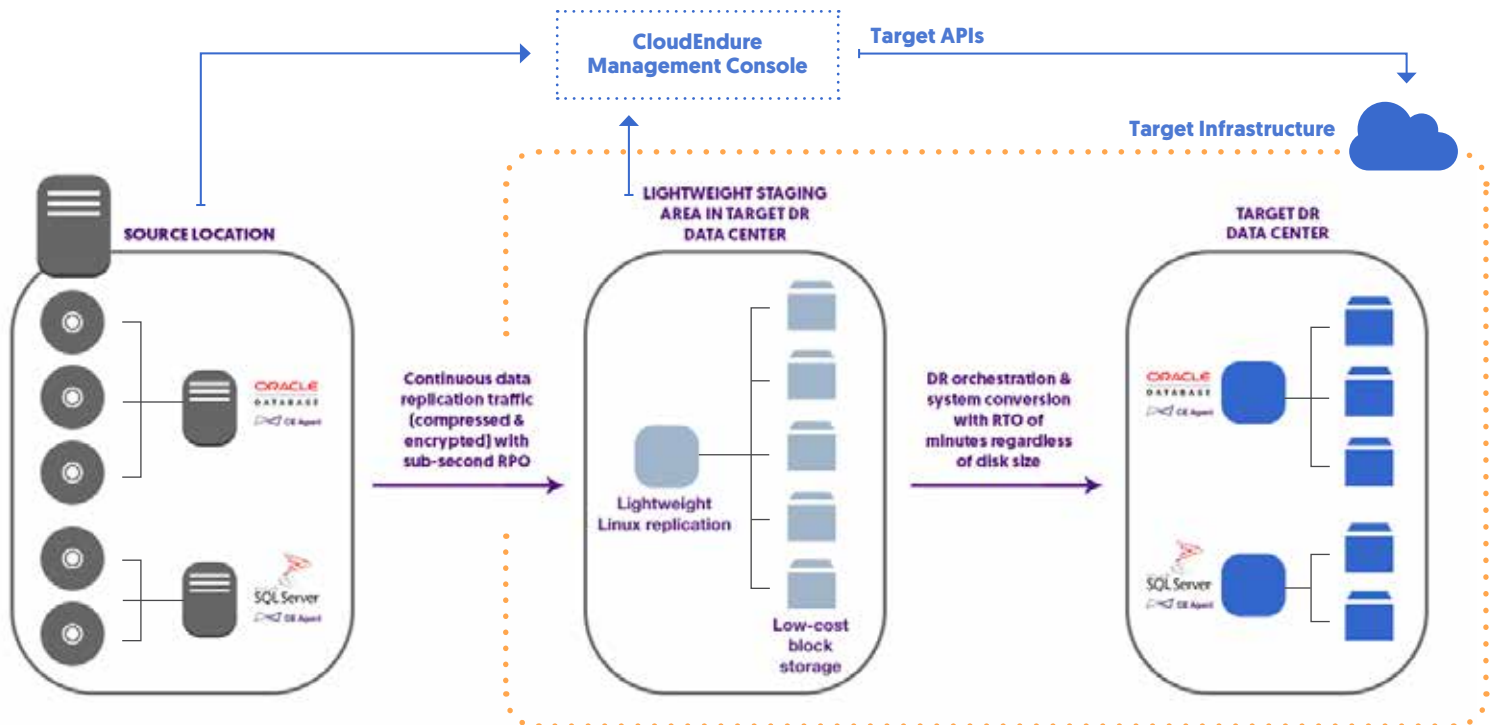
In order to ensure the highest level of security standards, data replication occurs directly from the customer's source infrastructure to the customer's target infrastructure, and can be restricted to private networks for better security, speed, and control. Data never goes through CloudEndure during the entire process. CloudEndure is ISO 27001 compliant and provides in-transit data encryption using AES 256-bit and supports at-rest data encryption in the customer's target site.



◀ **CloudEndure Console**
Point-in-Time Recovery

Point-in-Time recovery is critical in cases of database corruptions, accidental system changes, ransomware, or other malicious attacks on your environment.

Architecture of CloudEndure Technology



Each replication server can support a large number of source machines, significantly reducing compute costs for disaster recovery purposes. This is in contrast to traditional disaster recovery solutions, which require a constantly running target server for each source machine.