

How CloudEndure Disaster Recovery Works

THE TECHNOLOGY BEHIND CLOUDENDURE'S ENTERPRISE-GRADE DISASTER RECOVERY SOLUTION

Introduction

CloudEndure Disaster Recovery is a Software-as-a-Service (SaaS) solution that replicates any workload 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 of a customer's Disaster Recovery projects, CloudEndure provides APIs that enable developers to implement large-scale automation and other advanced capabilities.

CloudEndure Disaster Recovery is part of a comprehensive IT Resilience Suite, which includes Tier 1 Disaster Recovery, Tier 2 Disaster Recovery, and Continuous Backup. Customers can choose the most cost-effective option for each of their workloads based on criticality level and recovery requirements. This technical white paper primarily focuses on CloudEndure's Tier 1 Disaster Recovery solution.

Benefits of CloudEndure Disaster Recovery

CloudEndure Disaster Recovery enables organizations to quickly and easily **shift their disaster recovery strategy to public clouds, private clouds, or existing VMware-based data centers**. 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 infrastructure into **AWS, Azure, Google Cloud Platform (GCP), VMware, and OpenStack-based infrastructure**.

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
- ➔ Multiple IT resilience options, ensuring a cost-effective strategy
- ➔ 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

Continuous Data Replication

At the core of CloudEndure's technology is our proprietary **Continuous Data Replication** engine, which provides real-time, asynchronous, block-level replication.

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

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, security groups, replication tags, and more, through the self-service, web-based CloudEndure Console.

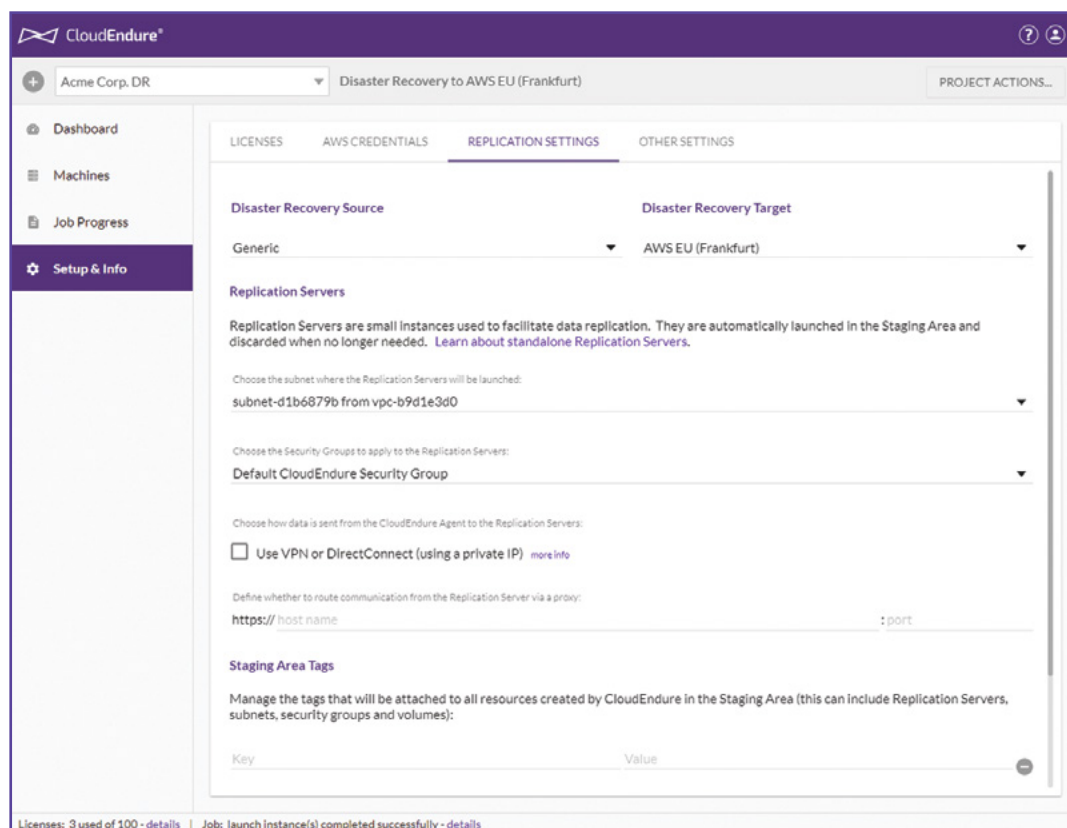
The initial replication can take anywhere from several minutes to several days, depending on the amount of data to be replicated and the bandwidth available between the source and target infrastructure. No reboot is required nor is there any system disruption throughout the initial replication.


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**.

CloudEndure Console Replication Settings

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





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 fallbacks right from the Console.

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 on the next page.

Low-Cost “Staging Area” in Target Infrastructure

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.

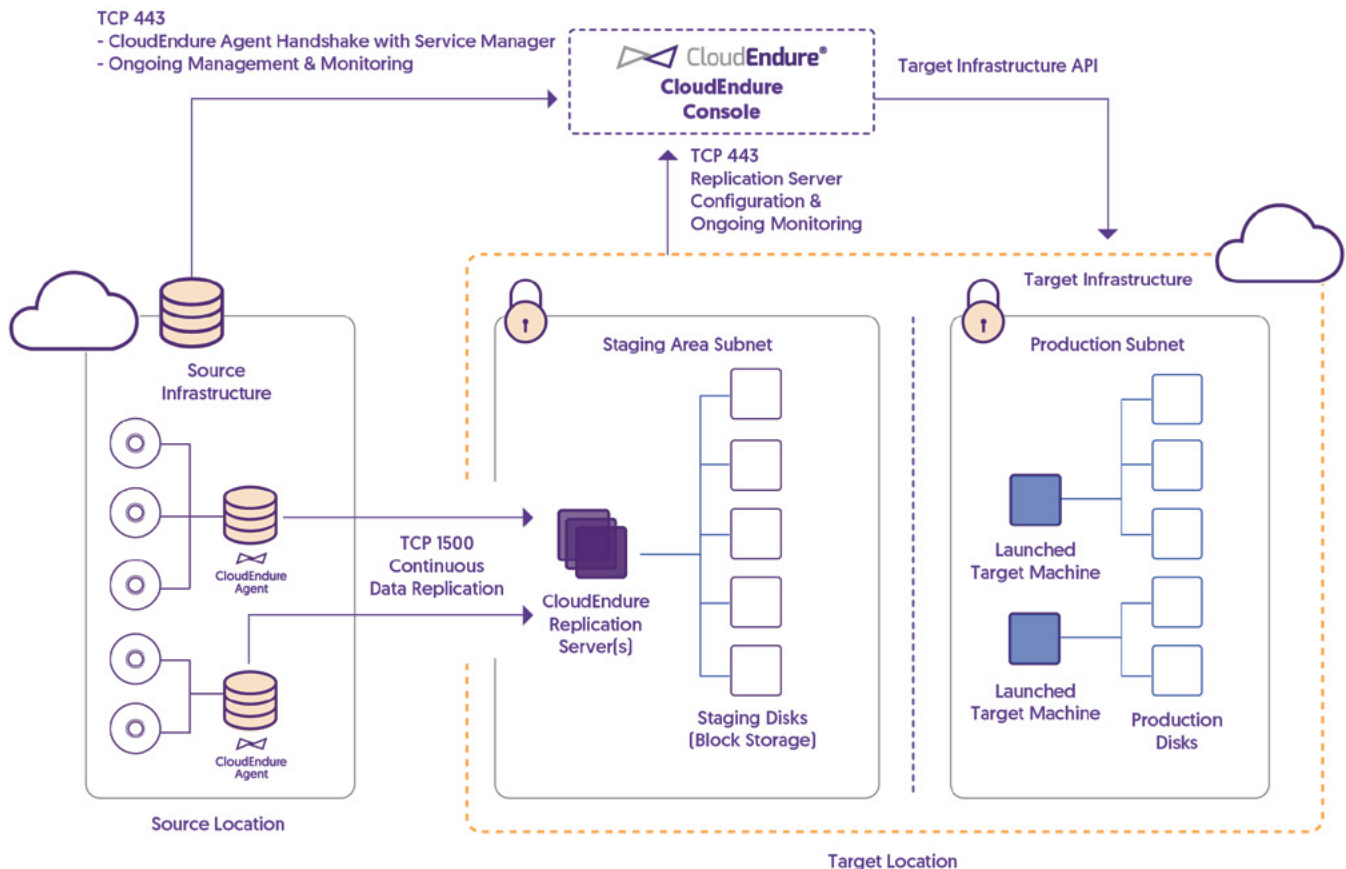
Whereas traditional disaster recovery solutions require duplicate hardware, compute, storage, networking, and software licenses in order to ensure data integrity and near-zero RPO, **CloudEndure’s “staging area” eliminates the need for duplicate provisioning of resources and reduces the TCO for disaster 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.

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

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.

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.

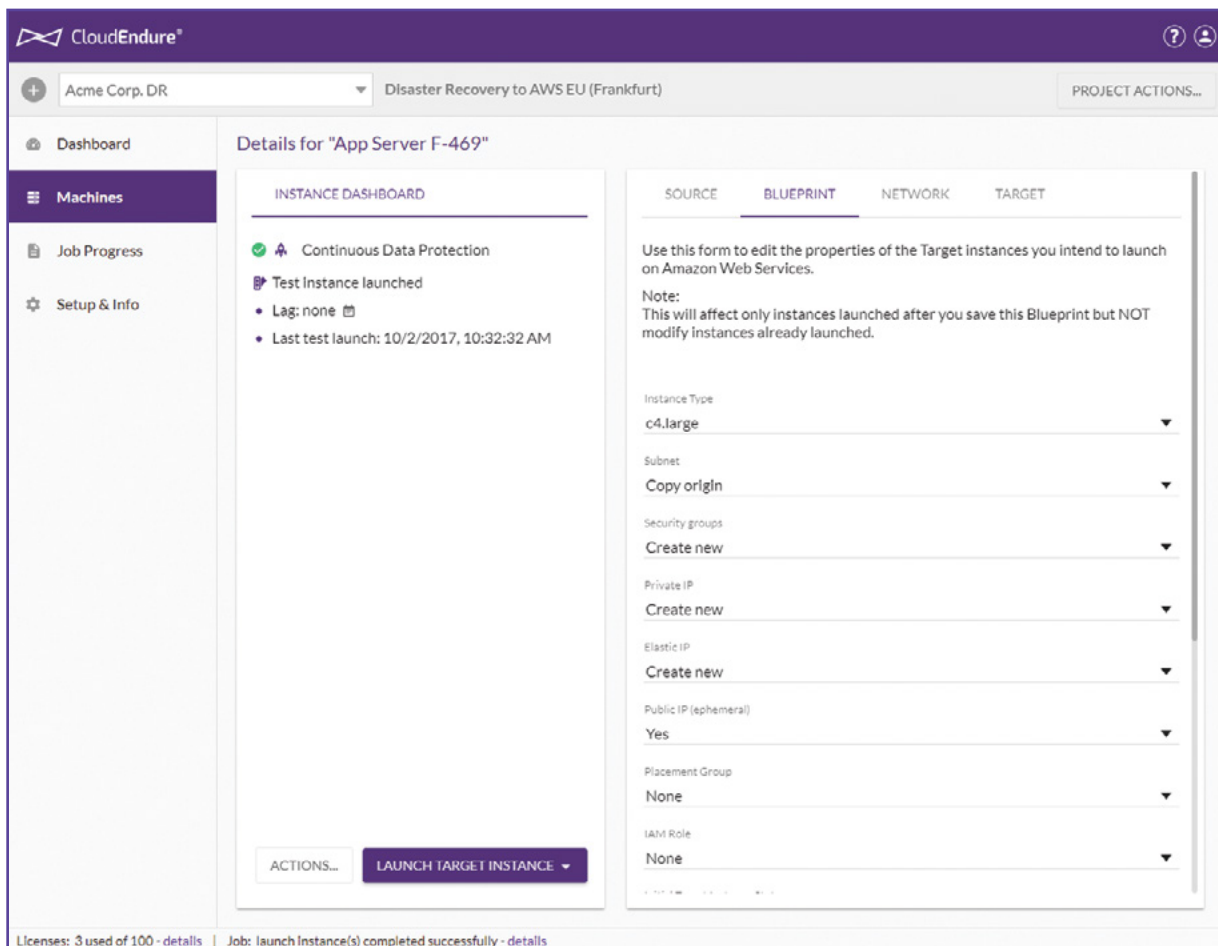
Automated Orchestration of Target Environment

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. Automated orchestration combined with machine conversion (detailed below) enable customers to achieve **Recovery Time Objectives (RTOs) of minutes**.

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



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, security groups, elastic IP, and more.

Automated Machine Conversion

When replicating machines across similar infrastructure, the replicated machines can boot natively in the target environment, as there are no significant differences in infrastructure. However, when replicating machines across dissimilar infrastructure, 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.

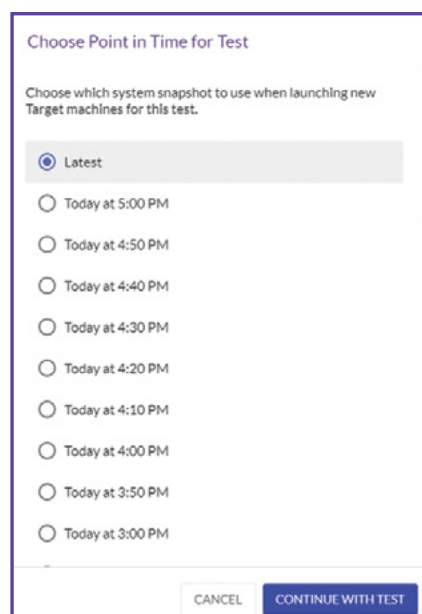
Multi-Tier Disaster Recovery and Continuous Backup

CloudEndure's IT Resilience Suite includes three solutions: **Tier 1 Disaster Recovery, Tier 2 Disaster Recovery, and Continuous Backup.** By offering three distinct solutions, CloudEndure enables customers to build the most cost-effective IT resilience strategy.

Tier 1 Disaster Recovery provides enterprise-grade protection for mission-critical workloads that require the most aggressive RPOs and RTOs. Tier 2 Disaster Recovery is used for business-critical workloads that can tolerate higher RPOs and RTOs. Continuous Backup is used to protect non-critical workloads. As a backup solution, it provides real-time server replication into low-cost object storage in the cloud or VMware, so that data can be restored to a preferred location.

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



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.

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.**

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

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.