

How CloudEndure Disaster Recovery Works

The Technology Behind CloudEndure's Enterprise-Grade Disaster Recovery Solution

Introduction

CloudEndure Disaster Recovery into Amazon Web Services (AWS) is a Software-as-a-Service (SaaS) solution. The solution is powered by innovative workload mobility technology, which continuously replicates applications from physical, virtual, or cloud-based infrastructure to a low-cost “staging area” (detailed below) that is automatically provisioned in any target AWS Region of the customer's choice. During failover or testing, an up-to-date copy of applications can be spun up on demand and be fully functioning in minutes.

Enterprises use CloudEndure to replicate their most critical databases, including Microsoft SQL Server, Oracle, and MySQL, as well as enterprise applications such as SAP. CloudEndure Disaster Recovery enables rapid recovery of the application, database, files, OS configuration, and system state — meaning that operations continue smoothly with fully functioning workloads. In addition to a self-service, web-based Console with centralized management for all of a customer's projects, CloudEndure provides APIs that enable developers to implement large-scale automation and other advanced capabilities.

Benefits of CloudEndure Disaster Recovery

CloudEndure Disaster Recovery enables organizations to quickly and easily shift their disaster recovery strategy to AWS from existing physical or virtual data centers, private clouds, or other public clouds, in addition to supporting cross-region disaster recovery in AWS.

CloudEndure Disaster Recovery 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 achieve sub-second Recovery Point Objectives (RPOs).

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

Benefits of CloudEndure Disaster Recovery include:

- ➔ Significant reduction in total cost of ownership (TCO) compared to traditional disaster recovery solutions
- ➔ Sub-second Recovery Point Objectives (RPOs)
- ➔ Recovery Time Objectives (RTOs) of minutes
- ➔ Support for any source infrastructure and any workload, including critical databases such as Microsoft SQL Server, Oracle, and MySQL, and enterprise applications such as SAP
- ➔ Automated failover to target AWS Region during a disaster
- ➔ Point-in-time recovery, enabling failover to earlier versions of replicated servers in cases of data corruption, ransomware, or other malicious attacks

- ➔ One-click failback, restoring operations to source servers automatically
- ➔ Unlimited, non-disruptive disaster recovery drills

- ➔ Replication is also supported between Regions or Availability Zones in AWS

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 the following types of source infrastructure:

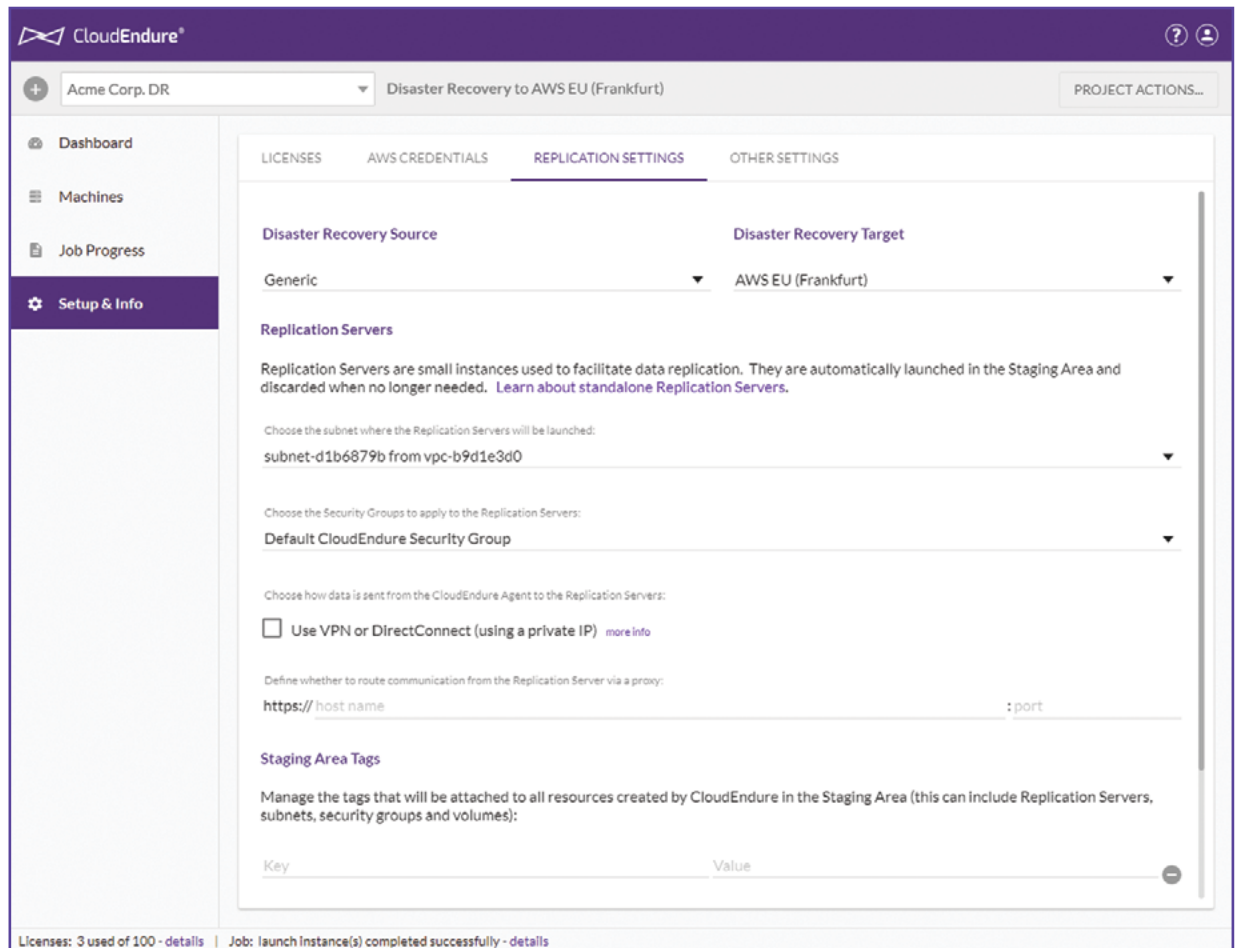
- ➔ Physical machines, including both on-premises and co-location data centers
- ➔ Virtual machines from any hypervisor, including VMware, Microsoft Hyper-V, and others
- ➔ Cloud-based machines, including AWS, Microsoft Azure,

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” that is automatically provisioned in a customer’s AWS account, in a target network of their choice. Customers define replication settings, such as subnets, security groups, and replication tags, 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 infrastructure and the target AWS Region. No reboot is required nor is there system disruption throughout the initial replication. After the initial replication is complete, the source machines are

CloudEndure Console Replication Settings

Customers select source infrastructure and target AWS Region, subnets, security groups, and more from the CloudEndure Console before beginning initial replication.



continuously monitored to ensure constant synchronization, up to the last second. Any changes to source machines are asynchronously replicated in real time into the AWS “staging area.”

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.**

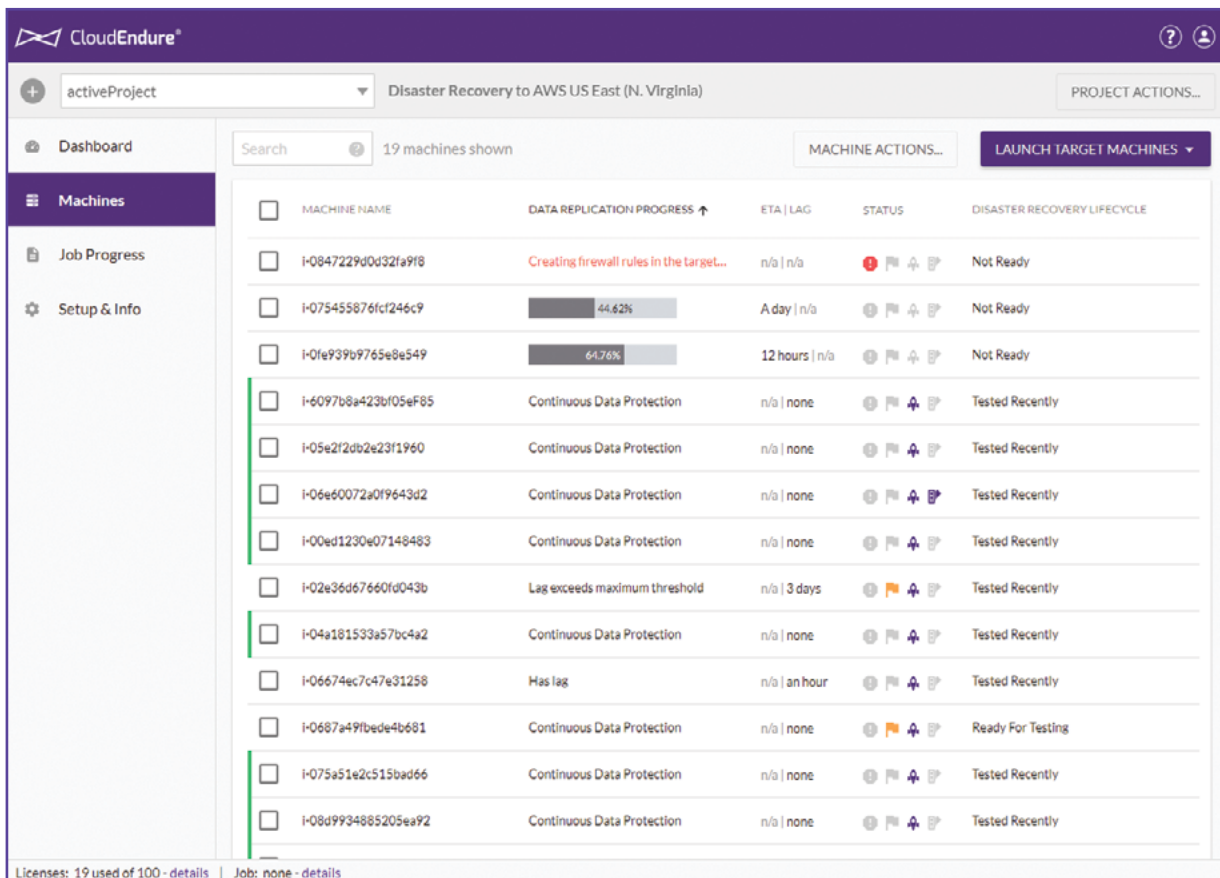
Continuous Data Replication allows customers to achieve sub-second Recovery Point Objectives (RPOs), as the data is always up-to-date and ready to be spun up as soon as a disaster strikes.

CloudEndure’s replication engine is application-agnostic, supporting all application types, including databases and legacy applications, any source infrastructure, and a wide array of operating systems.

Low-Cost “Staging Area” in Target AWS Region

CloudEndure maintains ongoing replication of source machines into a low-cost “staging area” in the customer’s preferred target AWS Region. The “staging area” contains cost-effective resources automatically provisioned and managed by CloudEndure to receive the replicated data without incurring significant costs. These resources include a small number of lightweight 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 sub-second RPO, CloudEndure’s “staging area” eliminates the need for duplicate provisioning of resources and thus **significantly reduces the TCO for disaster recovery.** The fully provisioned recovery environment, with the right-sized compute and higher



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.

performance storage required for recovered workloads, is only launched during a disaster or a DR drill.

Automated Orchestration of Target AWS Environment

CloudEndure’s orchestration engine **automatically launches fully operational workloads in the target AWS Region** at the time of failover or test. This automated process includes cloning disks from the “staging area” to target networks and provisioning all additional resources such as VMs, network interfaces, and firewalls.

The entire orchestration process typically takes minutes, depending on the general time it takes the machines to boot. CloudEndure’s orchestration process is not impacted by the number or size of volumes or the number of VMs.

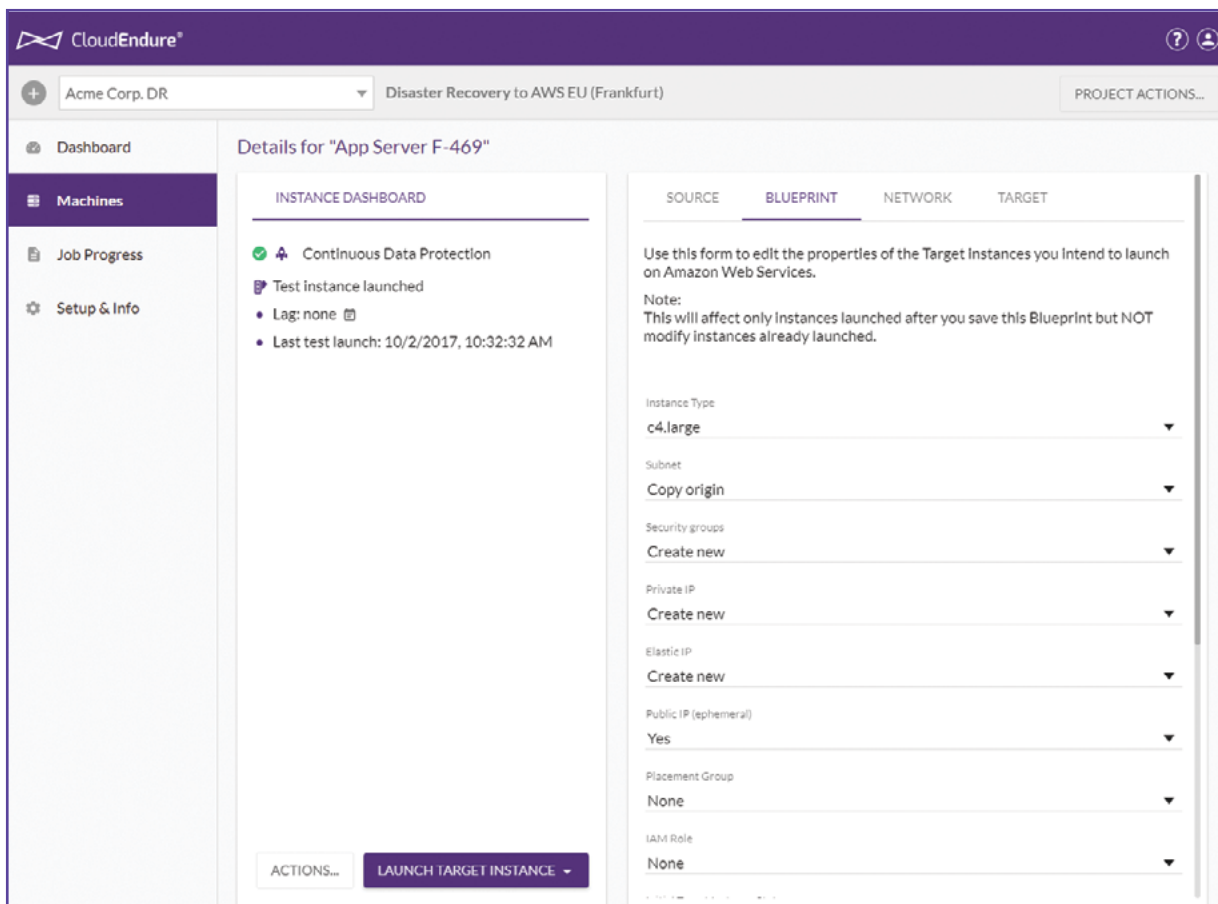
A large number of machines can be directly provisioned in parallel with a single click, using the appropriate predefined networking

and machine properties. Automated orchestration combined with machine conversion (detailed below) enables 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 AWS Region, instance/machine type, subnet, security groups, and elastic IP in the CloudEndure Console.

Automated Machine Conversion

When replicating machines across similar infrastructure, such as between AWS Regions or Availability Zones, 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, such as from on-premises or other clouds into AWS, machine conversions are required to ensure that the replicated machines can continue to run natively within AWS. This includes modifications to hypervisors, drivers, and other variations.



CloudEndure Console Blueprint Configuration Settings

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

Without proper conversion, such transitions between physical machines, hypervisor variations, or different clouds will result in non-bootable target machines.

CloudEndure addresses this using 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 Windows or Linux machines replicated from physical, virtual, and cloud-based infrastructure will natively boot and run transparently in the customer's preferred target AWS Region.

Choose Point in Time for Test

Choose which system snapshot to use when launching new Target machines for this test.

- Latest
- Today at 5:00 PM
- Today at 4:50 PM
- Today at 4:40 PM
- Today at 4:30 PM
- Today at 4:20 PM
- Today at 4:10 PM
- Today at 4:00 PM
- Today at 3:50 PM
- Today at 3:00 PM

◀ **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.

Applications*	Hardware Configurations*	Operating Systems*	Source Infrastructure*
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 (GCP)
Microsoft Active Directory	SAN-Based Disks	Oracle Linux	IBM Cloud
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, 2019	Oracle Cloud
Microsoft SQL Server		Windows XP, Vista, 7, 8, 8.1, 10	Physical Servers
MongoDB			VMware
MySQL			Xen
NGINX			
Oracle Database			
Oracle ERP			
Oracle Peoplesoft			
Pivotal CRM			
PostgreSQL			
Redis			
SAP CRM			
SAP ERP			
SAP S/4 Hana			

◀ **Applications, Hardware Configurations, Operating Systems, and Source Infrastructure 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.*

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, or other malicious attacks, customers may need to recover environments to previous consistent points in time. CloudEndure Disaster Recovery offers granular 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 utilizes the same agent-based block-level Continuous Data Replication engine, failback to source

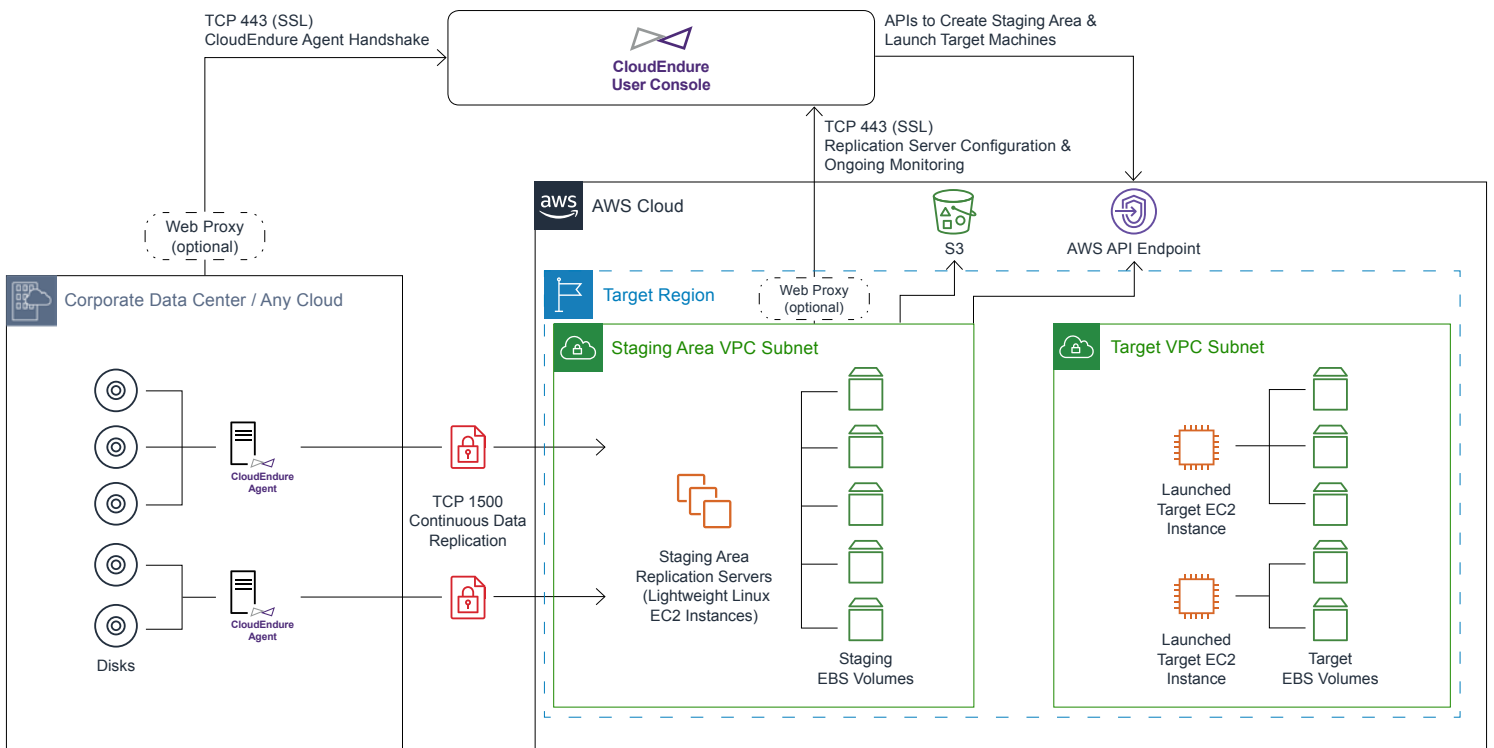
machines is rapid and **data is continuously synced in real time** from AWS until failback. CloudEndure's automated failback supports both full as well as incremental restores.

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 AWS Region, and can be restricted to private networks for better security, speed, and control. **Server 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 data-at-rest encryption in the customer's target AWS Region

CloudEndure Network Architecture



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.