CloudEndure®

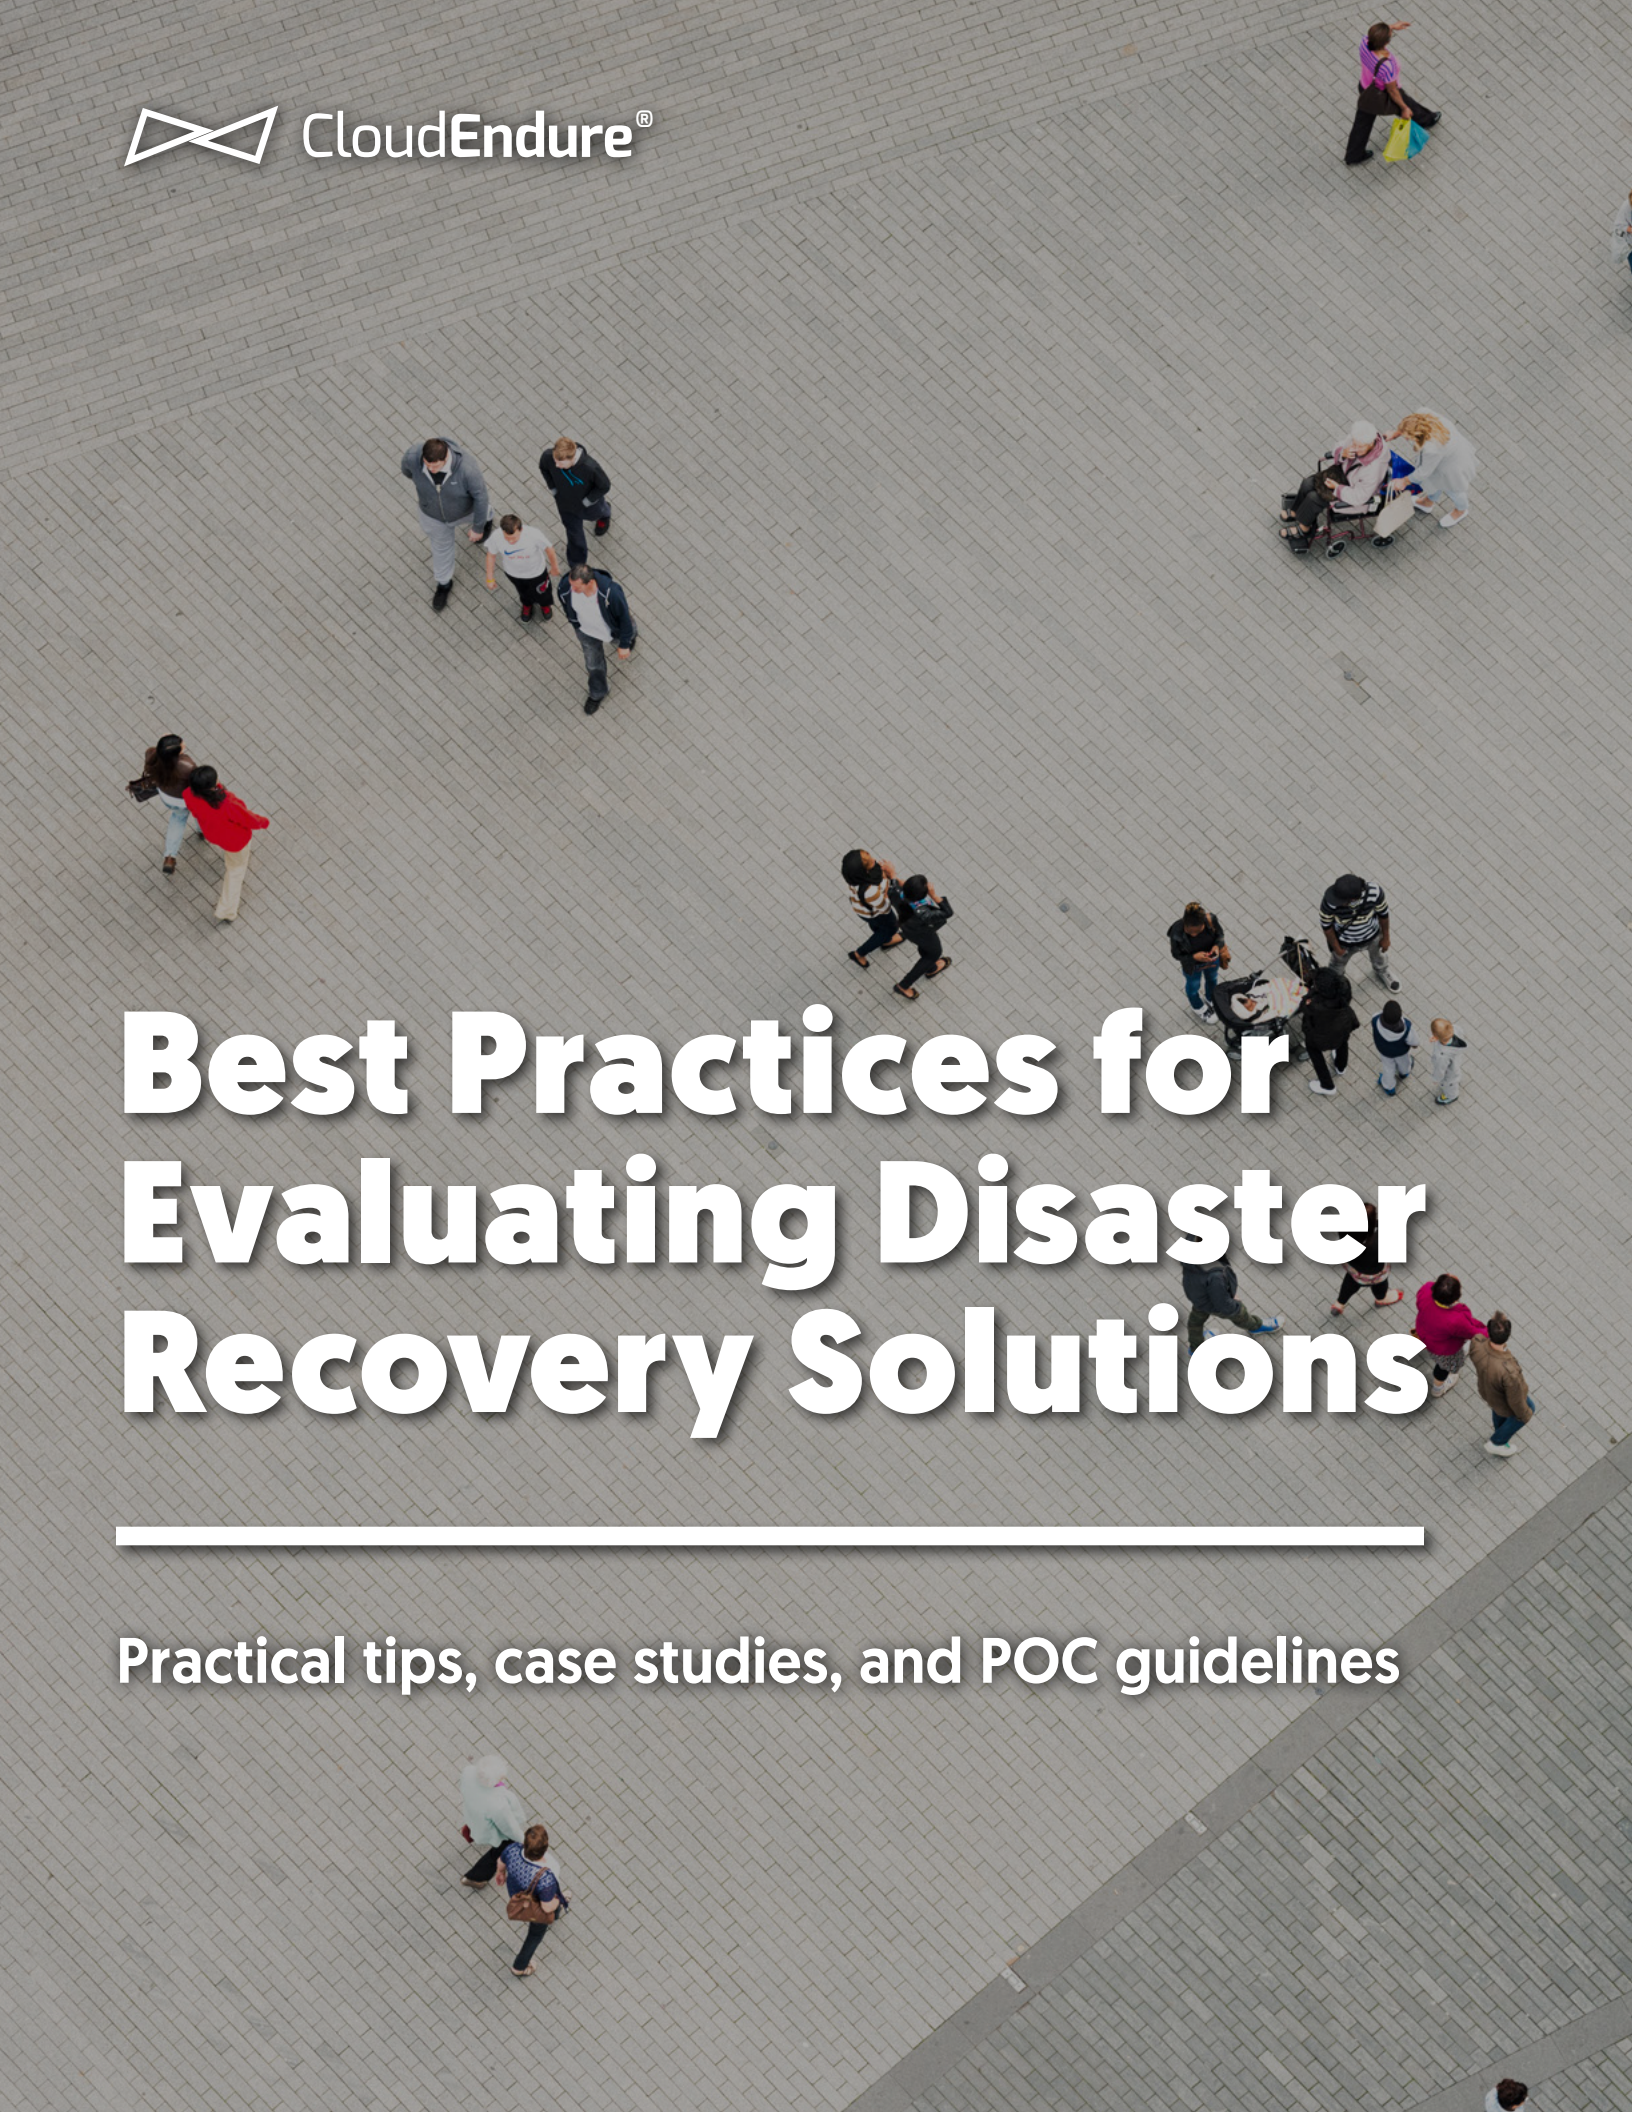

Best Practices for Evaluating Disaster Recovery Solutions

Practical tips, case studies, and POC guidelines

Recent advances in replication and recovery technology have dramatically changed the disaster recovery (DR) market. Businesses no longer need to settle for hours of downtime, nor do they need to pay exorbitant amounts for duplicate resources to ensure availability and data protection.

Based on our interviews with companies ranging from consumer product manufacturers to cutting-edge system integrators, it is clear that businesses are becoming more aware of the risks of a poor DR solution, and are taking a second look at their current setup. During their assessment, many companies find that their systems are expensive, resource-intensive, and difficult to scale — and don't necessarily provide an optimal level of protection from downtime and data loss.

The next logical step for these companies is to develop an alternative DR strategy. However, there is a wide range of DR solutions, each with varying levels of effectiveness. This white paper serves to guide organizations of all sizes with the key steps to follow when evaluating different DR solutions.

First Things First: Your List of Requirements

Before speaking with vendors, and in order to decide exactly which one of them is the best match, you should determine what exactly you need from a DR solution. What are your organization's biggest DR challenges? What are the critical features you can't compromise on?

1| Level of Protection

The first requirements you should determine are your Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs). In most cases, different servers will require different RTOs and RPOs, depending on whether they run mission-critical applications and if there are frequent data changes.

Most companies have servers that, if not operating, will have a direct impact on the business and cause financial losses. As Greg Ward, VP of Information Systems and Technology at Malibu Boats, a leading designer, manufacturer, and marketer of performance sports boats, explained, "If the technological systems go down, then our factory will shut down. If the factory shuts down, it's going to cost us a lot of money on a daily basis."

Because shorter RPOs and RTOs require higher costs to maintain, it is best to divide workloads into tiers. Mission-critical applications should be assigned to the highest tier and have near-zero RPOs and RTOs, meaning the need for a DR solution that can deliver real-time continuous data replication and spin up target servers in minutes. Mid-tier workloads require less aggressive RPOs and RTOs, so less sophisticated solutions should suffice. For workloads classified as low tier, such as those with infrequent data changes, a simple file-level backup system might be enough protection.

After grading their systems, many companies find that cloud-based DR solutions provide the best RPOs and RTOs and at the most reasonable cost. Although on-premise near-zero RPOs and RTOs are possible, achieving this is expensive and complicated.

2| Secondary Data Center vs. the Cloud

As companies grow, they often find that maintaining a secondary data center for DR is no longer a cost-effective strategy for long-term growth. Secondary data centers require duplicate resources for hardware, compute, storage, networking, and software licenses in order to ensure data integrity, so expenses can increase significantly over time.

CGS, a global business applications, enterprise learning, and outsourcing services company, recently decided to decommission their secondary facility and shift their DR strategy to AWS.

Companies like CGS that require enterprise-grade protection are finding that there are several key advantages to using cloud services for DR, which is now available at a dramatically lower cost than ever before:

- **Lower expenditure** – A good DR solution provides advanced replication technologies that leverage cloud infrastructure in a cost-effective manner. They enable a lower total cost of ownership because no expensive secondary DR site is necessary. In addition, with the right DR solution, you only pay for the cloud resources that you utilize during actual disasters and occasional testing.

  According to Michael Brandi, VP of the Technology Solutions Division at CGS, "when the software costs... are compared to CGS's current DR solution, CGS can save 50% of its DR costs annually with the new cloud-based disaster recovery solution."

- **Scalability** – Expanding companies that purchase a DR site often find that their DR setup cannot provide enough protection, storage, or capacity. In contrast, cloud-based DR solutions can scale to accommodate any size of business and deliver enterprise-grade agnostic protection for any application or database, without an impact on servers.

- **Performance & Speed** – For some companies, using the cloud for DR is a first step in their journey to the cloud. Once they experience the public cloud, many often decide to migrate their entire production environment.

  Malibu Boats is one such company. During their first disaster after shifting their DR strategy to the cloud, Malibu Boats managed to quickly failback to AWS and had everything back up and running within minutes. According to Ward, his IT team "quickly noticed that the performance

of the server on the AWS environment was at least twice as fast as it was on our current data center environment, using the same hardware specifications, the same memory, and CPU, which was quite eye-opening and helped us to move forward our plans to a fuller presence in AWS and the cloud."

3| Security and Compliance

Does your company have external security standards or compliance requirements to meet? Health Quest, an integrated delivery healthcare system, with more than 6,000 employees, and which services four hospitals, had to meet HIPAA (The U.S. Health Insurance Portability and Accountability Act) regulations for covering all services with a robust DR plan. According to CTO Rob Gilliland, Health Quest found that cloud-based services were the best way "to ensure our systems are replicated properly and that the environment is healthy."

In order to meet many compliance requirements, and to ensure the highest level of security standards, data replication should occur directly from the customer's source infrastructure to the customer's target infrastructure. It should also be restricted to private networks for better security, speed, and control. Moreover, most companies require in-transit and at-rest data encryption.

4| Features & Operations

Each organization has its own DR challenges based on their various types of storage, operating systems, and infrastructure. It is therefore imperative to begin the evaluation with a full picture of your IT environment.

Many organizations don't have the human resources or technical capabilities to handle DR with efficiency. For example, when an organization experiences rapid growth, coordinating various technologies to keep everything in sync and testing properly can became a real burden, and in some ways, unreliable. Such organizations should consider DR solutions that are easy to use and don't require ongoing maintenance.

Companies may also want to consider a Disaster Recovery as a Service (DRaaS) solution, which is fully managed by a

partner who ensures that systems are replicated properly and that the environment is healthy, and who administers failover support in the event of disaster. The quality of this kind of DR service depends on the particular DRaaS provider's technology, processes, and service-level agreements (SLAs).

Evaluating the Options in the Smartest Way Possible

After contemplating the above issues, many organizations decide that a cloud-based solution is their best option. The next step is to evaluate various solution providers. Specifically, organizations need to check public cloud platforms, such as AWS, Google Cloud Platform, and Microsoft Azure, and DR solutions that can integrate into the preferred cloud.

A common procedure is to review online data, form a short list, and then request demos, presentations, and answers to remaining questions. Research might raise some red flags, such as the need for multiple, time-consuming steps to set up and maintain the system; the difference between actual RPO and RTO and the expected objectives; or the need to pay for cloud compute fees for standby servers. (A quality DR solution will provide a low-cost, dormant staging area for this purpose, and require paid networking into the cloud only when disaster strikes or when performing a DR drill.)

It is also a good idea to speak with actual users of the software to get an objective, experienced opinion, and to obtain advice about the remaining evaluation steps.

It's All in the Proof — of Concept

Once you have narrowed your short list even more, it is time for a Proof of Concept (POC). A recommended procedure for a POC is to implement several scenarios with a selection of critical apps, representing the entire production environment. Testing should cover various critical system components including databases, a mix of physical and virtual machines, and (if relevant), large disk sizes. As explained by Michael Brandi at CGS, "Most vendors will say they have a solution, but it's often not as elegant as you would expect. Seeing how the process actually occurs in a demo or POC is critical."

A good POC will prove that critical applications are replicated and recovered quickly into the cloud without disruption or data loss, and demonstrate the solution's ability to seamlessly failback replicated servers to the source environment. It is also important to check if there is any load on servers due to replication. Of course, if the organization has set RPOs and RTOs, the POC should test these objectives.

Moreover, some DR solutions enable an easy POC setup through a guest account, as well as self-deploying agents. It is also helpful to receive fast response times during the POC process in order to allow the organization to make a fast decision, and which indicates the quality of service you can expect from the vendor.

Don't Forget to Test

Once a DR solution has been chosen, companies should conduct a comprehensive DR drill as soon as everything has been set up. Moving forward, DR drills should be conducted on a regular basis. It helps to have a solution that provides a single button for test activation as well as a one pane of glass interface, which allows the user to log in and see everything in real time, promotes a fluid testing process, and keeps track of the entire environment.

Regular testing might reveal gaps during initial failover from a network configuration, check proper handling of IP addresses, certain firewall rules, bandwidth requirements, and communication.

In addition, some organizations perform a small-scale failover test whenever they complete replication for any new workloads.

FAQs on Cloud-Based Disaster Recovery Solutions

Throughout our vast experience in working with partners and customers on their DR projects, as well as speaking to prospects, we hear many of the same questions and concerns about cloud-based DR solutions. We have chosen some of the most common questions and provided answers to them below.

1. How do I deal with the fact that my source infrastructure may be incompatible with the cloud infrastructure?

A good cloud DR solution would provide automated machine conversion, which ensures that any Windows or Linux machine coming from any source infrastructure (physical/virtual/cloud) will natively boot and run transparently in your preferred target. This will ensure your IT team doesn't have to spend days or weeks converting machines to make them compatible with your target cloud. It will also prevent human errors that often occur when this process is done manually.

2. What if I have legacy applications that were not built to be compatible with the cloud?

When replication is conducted at the block level, any file system or application can be transparently supported. Common workloads include the suite of databases and applications from vendors such as Oracle, SAP, and Microsoft, as well as proprietary legacy applications.

3. What are the average expected TCO savings for cloud-based DR vs. traditional DR?

From what we have seen, companies that shift their DR strategy to the cloud and decommission their secondary data centers save an average of 80% on DR total cost of ownership.

4. I'm in a regulated industry. Is cloud-based DR as secure as traditional DR, and would it meet the regulations I'm required to comply with?

The big public cloud providers are all compliant with most regulation standards. With regard to cloud-based DR solutions, many vendors offer complementary layers such as encryption at rest and in transit, the ability to replicate over a VPN, and other security and regulation features.

5. How is software licensing handled with cloud-based DR?

One of the easily overlooked but quite significant cost savings factors of using the cloud in conjunction with an advanced cloud-based DR solution is eliminating the need to purchase duplicate software licenses for your standby DR site. The most cost-effective DR solutions available today are able to maintain servers in real-time sync in a dormant "staging area" that is not running any licensed OS or applications. Only in the event of a disaster or a DR test, when servers are actually launched, are third-party OS and application licenses utilized.

6. Can cloud-based DR offer the same aggressive RPO and RTO that I can get with traditional DR solutions?

Definitely. In fact, with the right DR solution, you can get even better results for RPOs and RTOs when using the cloud. The only difference is that it will be much less expensive to achieve these aggressive recovery objectives.

7. How can I conduct DR drills with a cloud DR solution?

DR drills are much easier when using the cloud. When using an on-premise or DRaaS DR strategy, you have to ensure that the resources needed for the drill are provisioned and paid for in advance. In some cases, initiating a DR drill requires disrupting your source applications to avoid network conflicts.

When using the cloud for DR, you can simply request the resources when needed, and only pay for them upon usage. Furthermore, you can spin up your DR servers easily in complete isolation, thereby performing DR drills without any impact or conflict with your source applications.

8. If I run my DR servers in the cloud, once the disaster is over, how time-consuming and costly will the failback be to my on-premise facility?

With some solutions, this can be a cumbersome manual process of setting up your source servers and applications from scratch, moving the data, and then keeping it in sync until the point of failback. Other solutions allow you to simply reverse the replication back to your on-premise site (and keep the data in real-time sync) until you're ready to flip the switch and failback within minutes.

Why Do So Many Enterprises Turn to CloudEndure?

When evaluating DR solutions, many organizations discover CloudEndure and realize that it is the right fit for their needs. Why are so many enterprises choosing CloudEndure? By employing the following features, CloudEndure has established itself as one of the industry's top names in disaster recovery:

- **OS-level replication** enables support of any type of source infrastructure, including physical machines, virtual machines, and cloud-based machines, while there are no disk size limitations.
- **Continuous data replication** provides real-time, asynchronous, block-level replication, which means sub-second RPOs.
- **Low-cost "staging area"** in the target cloud contains cost-effective, cloud-based resources to continually receive replicated data, without incurring any significant costs. 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 machine conversion** ensures that any Windows/Linux machine coming from any source (physical/virtual/cloud infrastructure) will natively boot and run transparently in the customer's preferred target.
- **Point-in-time recovery** protects and recovers data and IT environments that have been corrupted in cases of database failures, accidental system changes, ransomware, and other malicious attacks.
- **Automated cloud orchestration** launches completely functional workloads in the target environment of the customer's choice. For disaster recovery purposes, automated orchestration, combined with machine conversion, enable customers to achieve RTOs of minutes.
- **Automated failback** utilizes continuous data replication to ensure rapid failback to source machines without losing data. CloudEndure's automated failback supports both incremental and bare-metal restores.

Architecture of CloudEndure Cloud Mobility Technology

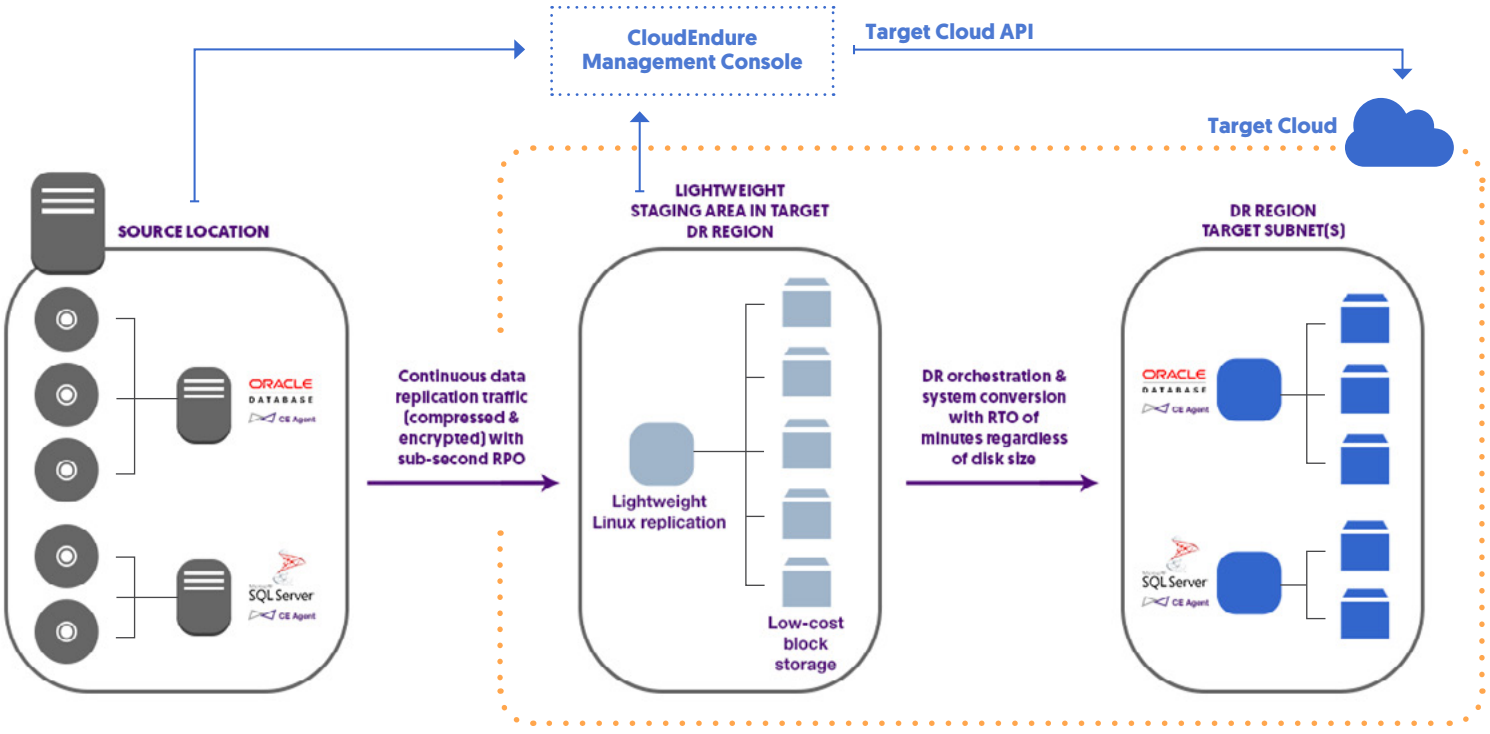

Each replication server can support a large number of source machines, significantly reducing compute costs for disaster recovery purposes. This is in contrast to traditional disaster recovery solutions, which require a constantly running target server for each source machine.