



Moving To Disaster-Recovery-as-a-Service (DRaaS)

A Guide to Selecting the Right Cloud Service

With more organizations recognizing the value of the cloud, the Disaster-Recovery-as-a-Service (DRaaS) market is expanding rapidly—by 2022, it's poised to grow into a [\\$12.5 billion market at an impressive](#) compound annual growth rate [of 42%](#). This growth has also resulted in an increasing number of players in this market, leaving organizations with a big question: how to choose among a vast sea of DRaaS providers? In this guide, we'll walk you through six critical questions to ask and features to look for when making that important decision to partner—a decision that will enable your organization to survive through disastrous events and come out on the other side with limited impact. Let's get started.



1 How quickly will I recover after a disaster?

When it comes to disaster recovery (DR), every hour that your business spends offline costs you—in both revenue and reputational damage. While the [cost of downtime](#) varies based on size and industry, it's pegged at a [whopping \\$540,000/hour](#) for most businesses, and it can jump up to [\\$1,000,000/hour](#) for Fortune 1000 companies.

For this reason, it's important to understand the DRaaS vendor's average recovery point objective (RPO) and recovery time objective (RTO)—two important KPIs used to measure the effectiveness of a DR solution. RPO is the time interval during which you will lose data (your “appetite” for data loss); RTO shows how quickly you can bring your production systems live again (the time that you can afford to wait without losing customers). For data-dependent organizations, the objective is to achieve a near-zero RPO and to minimize RTO as much as possible.

It's also worth asking about the approaches that your vendor is undertaking to meet these targets. Some vendors may only be replicating backups to the DR site and hydrating the data at regular intervals—this approach increases recovery time and is not a preferred option for organizations and businesses that are heavily data-dependent. On the other hand, a thorough DR solution will ensure asynchronous replication of data from the production systems to the backup sites, ensuring near-zero data loss.

At Enquizit, for instance, we use AWS CloudEndure Disaster Recovery (CEDR) as an essential component in our approach to a comprehensive continuous replication solution that supports multiple replication patterns including database and file. This model, compatible across diverse platforms, ensures near-zero data disruption and quick recovery using native AWS automation and orchestration capabilities.

RPO: the amount of time your systems can be down before negatively impacting your business.

RTO: the amount of time it takes to bring your production systems live again.

For data-dependent organizations, the objective is to achieve a near-zero RPO and to minimize RTO as much as possible.

In the case of Florida State University, Enquizit delivered near-zero data loss with less than 1 hour of RPO, a 4-hour RTO for Tier1 services, and an overall RTO of 12 hours.

One of our strategies is to preserve the production system environment attributes such as namespace, host interfaces, and logical network addressing in the DR Testing and DR Ready environments resulting in fewer differences between production and failover site. This approach enables applications to recover faster, reduces application testing complexity, and minimizes the gaps between RTO tiers (RTO tiers are a way of grouping systems and applications based on their business impact and promising service levels accordingly).

Finally, make sure your vendor is transparent about if, and how much, the RPO and RTO targets would be compromised when you scale up your infrastructure in the future. Your DRaaS provider's ability to maintain consistent performance with the ability to scale should be an important decision criterion for you.



② How long will it take to build the DR solution?

Building a customized DR solution for your organization can become a time-consuming exercise, even when a product is available off-the-shelf. There are many complexities involved in the transition, but let's look at two common problems:

1. Creating a Baseline


First, you need to baseline your existing infrastructure, inventorying every element that needs to be replicated.

2. Building a Design

Then, you'll also need to design the complete solution and figure out what other services might be required to build out the solution quickly.

As a result of such challenges, the time-to-market can stretch for a long duration—sometimes more than a year.

One of the best ways to expedite the deployment process is to work with integration partners who have technological depth and [prior experience in implementing cloud migrations in enterprises of different verticals and sizes](#). For example, our proprietary [SkyMap™ tool](#) is designed to use AWS artificial intelligence and machine learning to determine quicker data-driven paths for DR migration scenarios.



Florida State University's DR project was designed and launched by Enquizit in just 6 months – between September 2019 and March 2020.



Find out if your vendor adopts agile methodologies, like Scrum, and offers dashboards that can highlight problems early, as these capabilities will help reach your milestone faster.

With the integration of AWS services and automation capabilities, SkyMap™ can drastically reduce time-to-market as follows:

1. Automating baseline discovery

SkyMap™ helps orchestrate the discovery effort involved in creating a system of record of your organization's infrastructure.

2. Designing an end-to-end solution

SkyMap™ also helps architect a comprehensive DR solution with AWS CloudEndure at the center, leveraging other AWS services such as CloudFormation and Step Functions to complete the picture.

Additionally, your DRaaS provider needs to showcase superior project management skills, since migration to a cloud-based DR solution will require extensive communication and coordination, early identification of gaps, and continuous feedback. Find out if your vendor adopts agile methodologies, like Scrum, and offers dashboards that can highlight problems early, as these capabilities will help reach your milestone faster.

3 How many steps are in the runbook? Can I see a sample?

A DR runbook is a manual outlining the necessary steps to recover from an interruption or disaster. It offers instructions for personnel in the event of a disaster, including both infrastructure and process-related details. When you evaluate a DR solution, it's a good idea to look at the provider's sample runbook and assess if it's comprehensible and geared towards meeting your RTO target.

An effective runbook should emphasize two elements—simplicity and automation:

1. Simplicity

The number of steps to failover to the DR environment, during testing as well as during an actual event, should be as few as possible, translating into less overall effort required to meet a competitive RTO.

But the problem is that runbooks are frequently so outdated that when a DR exercise occurs, organizations need to do ad hoc remediation activities, such as patches and upgrades, to even get to the point of being able to run the exercise. Many times, DR sites are outmoded, with configurations not in sync with production data centers. That means that at the moment of failover, a lot of time is spent on 'production catch-up,' trying to fix the configuration drift between the two environments.

Bottom line? A simple and short runbook is one indicator that the DR environment is operating at a high level of readiness.



Enquizit's approach to implementing DRaaS focuses on minimizing manual intervention and greatly leveraging automation in the DR runbook.



Enquizit uses a plug-and-play approach with code; at least 80% of our automation scripts are client-agnostic and can be highly reused. This allows our clients to easily scale in the future.

2. Automation

For a runbook to get to the point of being simple and effective, the DR environment should be significantly automated and not dependent on manually intensive processes that tend to introduce errors and increase RTO.

For instance, CEDR offers built-in automated replication and orchestration services, including scripts that execute post-launch tasks on the AWS infrastructure. Regardless of whether you're executing this in a test environment or during an actual event, this level of automation reduces the configuration drift that can occur between production and DR and helps you get the alternate site ready in a matter of hours.

Additionally, as discussed in the previous section, the architecture takes advantage of other AWS services such as CloudFormation, Step Functions, and Lambda, which bring their own automation capabilities. [Lambda functions](#), for instance, have specific custom logic and are invoked from CloudFormation and the AWS console to help automate certain tasks.

4 How flexible is the DR solution? Will it work with my other technologies?

Moving to DRaaS is only convenient if the cloud solution is flexible and adaptable enough to integrate with some of your legacy systems. It's important to ask your vendor whether they have the capacity and flexibility to integrate with these other technologies that you are using.

For instance, many organizations and universities rely on Oracle PeopleSoft applications. But several well-known DR solutions do not offer the ability to integrate with Oracle. This can turn out to be a major roadblock during implementation. In addition, your integration partner needs to have expertise on both the application as well as on the cloud solution to ask the right questions (e.g. which AWS instances—AWS EC2 or AWS RDS—would be a better option for your particular use case when it comes to backing up the Oracle nodes?)

You may also want to integrate other enterprise technologies or vendor solutions that you currently use for security, operational management, and identity management into the architectural planning and design of the DR solution. Can your DR provider offer integration with configuration management platforms such as Ansible, Chef, and Puppet, or with directory services such as Active Directory? It's essential to consider the platforms that support your core architecture when you make the choice of a DR solution—because you need them to work together.

Enquizit supports multiple replication patterns—not just for applications and servers, but also for databases, network storage, and other enabling platforms. While many DR solutions offer only selective replication, Enquizit's end-to-end model ensures that no system, technology, or platform (legacy or recent) is left behind during disaster recovery.

For Florida State University's DR solution, Enquizit used a combination of replication patterns:



Host replication using CloudEndure DR



Oracle Exadata host platform with RAC replication to EC2 using Active Data Guard



File replication using SoftNAS

5 Once the cloud solution is in place, how easy is it to operate? Will my team be able to run it easily?

Maintenance and management of a DR solution can present challenging scenarios throughout its lifecycle. As your IT team is the first line of support at this stage, you need a DRaaS solution that's easier to manage when compared to on-premise data centers or other cloud solutions.

Let's look at how the AWS cloud tackles three critical areas of the DR lifecycle. These strategies ensure that disaster recovery is not an overwhelming process but remains hassle-free for your teams.

1. Maintaining a superior state of readiness

[The AWS platform supports many popular DR architectures](#) from “pilot light” environments that are ready to scale up at a moment's notice to “hot standby” environments that enable rapid failover. In a pilot light environment, the most critical elements of systems are already configured and running in AWS. At the time of recovery, you would rapidly provision a full-scale production environment around the critical core.

It's therefore important for your pilot light environment to support replication of all systems—physical and virtual servers, databases, networks, storage, and other peripheral systems. Blueprints, which are detailed templates that contain configuration details of how machines need to be provisioned, should be accurate and readily available.



Florida State University conducted a business impact analysis (BIA) that identified nearly 100 critical applications supported by Linux and Windows servers in the production data center that made up the scope of this next-generation DR solution. AWS has data centers in 18 regions around the world (5 in the United States).

The systems running in the pilot light environment need to be managed just like any other production system—patches and upgrades need to happen in sync with live environment, applications need to be continuously monitored, and issues need to be remediated with similar SLAs to that of production systems.

Your vendor should be able to help you here by finetuning your enterprise architecture management policies at the time of set-up of the DR environment (e.g. add test hosts to automated scripts that deploy patches).

2. Managing replication and governance

AWS CEDR has a separate staging VPC where the replication servers are hosted. CEDR Blueprints are stored in the CloudEndure portal and in S3 for safekeeping. The good news is that once it's in place, CEDR is largely self-managing, requiring little staff support in maintaining continuous data replication and protection.

To monitor issues that arise during replication, email notifications alerting on events such as replication lag can be

set up from the CE console. Additionally, customers can also use Enquizit's dashboard to view near real-time health stats of various CloudEndure elements in a single pane, rather than navigating different parts of the CE console.

The advantage of moving to AWS cloud is the availability of many AWS services that streamline the DR process. For instance, you can achieve improved operational and security governance using [AWS Control Tower](#).

3. Managing scalability

Adding a new CEDR host to the existing set-up is as simple as installing the CEDR agent on the server that you want to replicate. For Enquizit customers, our highly reusable parameterized scripts will help you add a new host with less effort and no modification to the runbook.

When you shortlist a DRaaS solution, verify that it offers a high level of readiness, replicates data continuously with hardly any manual intervention, and scales up easily. These features will aid your IT teams to perform their tasks with no stress.



6 How cost-efficient is your service?

Finally, we come to the cost analysis. The pay-as-you-go model of the cloud infrastructure is a big factor that works in your favor, as you do not have to lease on-premise data centers and can both avoid capital expenditures and reduce operational expenditures.

But how do you differentiate between the pricing structures of different DRaaS providers? Here, it's important to remember that you cannot just compare price tags; you also need to weigh the costs against your requirements and the RPO/RTO targets that you are looking to achieve. Performing a business impact analysis to evaluate if the cloud solution meets your expectations should also include the five questions that we have already discussed above.

Enquizit employs AWS CEDR in the following ways to keep operational costs low:



Offering replication and orchestration services that are focused on maintaining a low price point with up to 70% storage compression.



Allowing the option of running only essential systems and services (pilot light) without compromising performance.



Creating a temporary environment for DR tests, which is deployed and wiped on-demand, thus hugely reducing costs associated with mandated confidence testing.



Using the infrastructure-as-code methodology with an extreme focus on building automation scripts, thereby bringing down the costs of continuous replication by a significant margin.



Offering flexibility to scale up your infrastructure without much upfront cost.

When you are evaluating a DRaaS provider, these are some approaches to watch out for to decide if the provider is delivering a cost-effective solution.

Why SkyMap™ with CloudEndure is the Right Choice

Selecting the right DRaaS provider is a balance of three main factors: product capability, time, and cost. Capability refers to the features you get as part of a solution-set from an experienced DRaaS partner, along with how easy it is to set up, implement, maintain, and scale for future growth and change. Time refers not only to your RPO and RTO targets—how much data you can afford to lose and how long you can afford to be offline—but also to how quickly your solution can be implemented. The cost component touches each of the first two factors and should also be weighed against current costs of maintaining, upgrading, and paying for real estate for on-premises data centers.

Enquizit's comprehensive DR solution-set is designed to tick all these boxes. As an integration partner, we offer unique automation capabilities with SkyMap™ which, along with AWS CloudEndure, will make your DR transition to the cloud a seamless, and even enjoyable, experience. To learn more about our DRaaS architecture and solution, visit our product page at <https://enquizit.com/services/disaster-recovery/>.

