

DISASTER RECOVERY FOR THE CLOUD



In recent decades, [cloud computing](#) has gained popularity due to its range of benefits to business organizations:

- Cost optimization
- Access to high performance IT infrastructure
- [Security](#) and compliance
- Ease of doing business

However, these advantages are only realized as long as the service is available and functioning as per the expected reliability standards.

In order to maximize the [reliability of IT services](#) delivered from off-site cloud data centers, cloud vendors and customers (should) follow Disaster Recovery strategies. These practices are designed to mitigate the risks associated with operating mission-critical apps and data from cloud data centers, that are not immune to:

- Natural disasters
- Cyber-attacks
- Power outages
- Networking issues
- Other technical or business challenges affecting service availability to end users

Unplanned downtime cost businesses over \$80,000 per hour in data center downtime according to

a recent research. While large enterprises may be able to contain the financial damages associated with downtime incidents, small and midsize businesses experience the most damaging consequences. Research suggests that organizations without an adequate disaster recovery plan go into liquidation within 18 months of suffering a major downtime incident.

This makes disaster recovery planning critical to business success amid growing dependence on cloud-enabled IT services, [cybersecurity](#) issues and power outage concerns.

What is disaster recovery?

[Disaster recovery \(DR\)](#) is a component of security planning that constitutes the technologies, practices, and policies to recover from a disaster that impacts the availability, functionality and performance of an IT service. The disaster may result from a human, technology or natural incident.

Disaster recovery is a subset of business continuity that deals with the large picture of avoiding disasters in the first place. While business continuity involves the processes and strategy to ensure a functioning IT service during and after a disaster, the component of disaster recovery involves the measures and mechanism that help regain application functionality and access to data following a disaster incident.

The following is a brief guide to get you started with your disaster recovery planning initiatives.

Planning & preparation

Disaster recovery planning is unique to every organization and depends on the metrics that are best considered to evaluate the recovery of an IT service following a disaster. Organizations need to identify the resilience level for their development, testing, and production environments, and implement disaster recovery plans accordingly.

The metrics in consideration could include:

- Recovery Point Objective (RPO), the age limit of business information that must be recovered since the disaster
- Recovery Time Objective (RTO), the acceptable time for recovery during which the IT service remains unavailable

These metrics should be aligned with the organizational goals of business continuity and must evolve over time as you scale and face different challenges in achieving these goals.

For customers of [cloud infrastructure services](#), the requirements on these metrics should be defined in the [SLA agreement](#). High availability architecture such as hybrid and multi-cloud environments offer improved operational performance in terms of service availability. However, you must consider the investment option and tradeoff between cost, availability, performance, and other associated parameters.

Disaster recovery best practices

The following best practices should be employed in developing a disaster recovery program for your organization:

- **Define disaster.** Understand how your organization defines a disaster.
- **Define your requirements.** Understand your RPO and RTO requirements for different workloads and applications.
- **Create a disaster recovery practice.** How do you re-evaluate disaster recovery on [an ongoing basis](#) to account for changing technical and business requirements?
- **Test your plans.** Is your organization capable of realizing a disaster recovery plan in real scenarios? Consider employee awareness and training, disaster recovery exercises, and drills.

Know your options

Disaster recovery solutions may involve a diverse range of options for different DR goals. A well-designed strategy focuses on an optimal tradeoff of cost investments, practicality, and IT burden, with the disaster recovery performance.

For instance, if a car risks a puncture during driving, would you rather:

- Use expensive run-flat tires?
- Use a regular tire and keep a spare tire with a replacement kit in the car?
- Use a regular tire, have no spare, and rely on roadside assistance to replace a flat tire?

Each option have their own set of implications and require a strategic assessment of the disaster recovery goals. It may be possible for organizations to follow a holistic disaster recovery plan that incorporates different disaster recovery patterns for different use cases as appropriate. For instance:

- A mission-critical app may require short RTO/RPO objective.
- An external marketing database may not impact business operations for long duration following a disaster.

Testing your disaster recovery capability

Organizations can develop the most applicable and appropriate disaster recovery program—yet fail to implement the measures in practical, real-world environments. These limitations are often caused by limited employee training and failure to account for real-world situations that may have been ignored during the disaster recovery planning stages.

The proof is therefore in the testing of your disaster recovery program at frequent and regular intervals. These intervals may range from 1-4 times per year, although some fast-growing organization may even resort to monthly testing exercises depending upon their technical requirements or regulatory concerns.

The testing procedures should extend beyond the technology capabilities and encompass the people and processes.

Disaster recovery simulations can help organizations understand how the technology will behave in transferring workloads across geographic locations if the primary data center is hit with a power outage. But what about the workforce responsible for executing the policies and procedures designed to streamline the disaster recovery process? This means that the disaster recovery program should also consider the education and training of employees responsible for executing key protocols to recover from a disaster situation.

Finally, it is important to keep up to date documentation on the disaster recovery performance

during exercises as well as real-world disaster incidents. Use this information as a feedback loop to tune your disaster recovery capabilities based on your organizational requirements.

Disaster recovery for different cloud architecture models may be treated according to the impact on business and the technical requirements. For instance, [multi-cloud environments](#) may be less prone to disaster situations as per the appropriate SLA agreements associated with multiple data center locations, RPO/RTOs and other metrics.

Therefore, you must evaluate which cloud service model optimally fulfills your disaster requirements on different apps and data sets used to perform daily business operations.

Related reading

- [BMC Multi-Cloud Blog](#)
- [Cloud Infrastructure Explained](#)
- [Cloud Monitoring: Choosing the Right Metrics](#)
- [Hybrid Cloud Governance & Compliance](#)