

***Disaster Recovery
Best Practices***

DISCLAIMER

This document has been prepared by Cloud Management Office (CMO) under the Ministry of Electronics and Information Technology (MeitY). This document is advisory in nature and aims to provide information in respect of the GI Cloud (MeghRaj) Initiative.

Certain commercial entities, technology, or materials may be identified in this document in order to describe a concept adequately. Such identification is not intended to imply recommendation or endorsement by MeitY.

While every care has been taken to ensure that the contents of this document are accurate and up to date, the readers are advised to exercise discretion and verify the precise current provisions of law and other applicable instructions from the original sources. It represents practices as on the date of issue of this document, which is subject to change without notice. The readers are responsible for making their own independent assessment of the information in this document.

In no event shall MeitY or its' contractors be liable for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information) arising out of the use of or inability to use this Document.

Table of Contents

1	Purpose.....	6
2	Background.....	7
3	Why use Cloud Disaster Recovery?.....	9
4	Disaster Recovery Principles.....	11
5	Types of Disasters.....	12
6	Guidelines/Best Practices for adoption of Disaster Recovery.....	14
6.1	Identify criticality of Data & Applications.....	15
6.2	Selection of DR Site Architecture.....	17
6.3	Selection of Replication Technology.....	18
6.4	Understanding Bandwidth Requirements.....	20
6.5	Disaster Recovery as a Service (DRaaS).....	21
6.6	Documenting DR Plan.....	23
6.6.1	Roles and Responsibility.....	23
6.6.2	Segregation of Responsibility between CSP, MSP and a User Department.....	24
6.6.3	Scope and Dependencies.....	25
6.6.4	Service Level Agreement.....	25
6.6.5	Validating DR Readiness.....	26
6.6.6	Hybrid Cloud DR Scenarios.....	27
6.6.7	Government Laws and Regulations on Disaster recovery Site.....	27
Annexure 1	29
	Terminologies related to DR.....	29
	Cost considerations for setup of Disaster Recovery*.....	30



1 Purpose

MeitY has introduced MeghRaj initiative to utilize and harness the benefits of Cloud Computing in order to accelerate delivery of e-services in the country. It has led to a significant adoption of Cloud technology across Government Departments. As there is a significant upsurge in digital information throughout the Government eco-system, there is an increased need of preparing our digital ecosystem to overcome any disasters, without a considerable impact on public services. This document will assist User Departments in evaluating and considering the best practices suitable for their respective departments in terms of Disaster recovery and ensuring business continuity.



Adoption of Disaster recovery setup is important for all Departments to maintain availability of Government Operations and resiliency of data/applications.



2 Background

As there is a constant rise in information systems and electronic data, the rise in vulnerability of such data has also increased exponentially. The disruptions can be seen ranging from mild (e.g., short-term power outage, disk drive failure) to severe (e.g., site destruction, fire). Though the vulnerabilities may be minimized or eliminated through management, operational, or technical controls, as part of the departments resiliency effort, however, it is virtually impossible to completely eliminate all risks.

One of the challenges for User Departments is ensuring the operations remain unaffected, even during adverse times. Disasters can strike at any moment, leading to socio-economic and reputational losses.

The guideline focuses on describing Disaster recovery planning and detailing out the considerations and best practices which should be followed to mitigate the risk of system and service unavailability by providing effective and efficient solutions to enhance business continuity.

What is Disaster Recovery?

Disaster Recovery (DR) aims at protecting the Department from the effects of significant catastrophic events. It allows the Departments to quickly resume mission-critical functions after a disaster. Below figure explains various possible disasters that can take place.



Figure 1: Disaster scenarios

The goal for any Department with DR is to continue operating as close to normal as possible. Preparing for a DR requires a comprehensive approach that



Disaster Recovery Best Practices in Cloud Computing

encompasses hardware and software, networking equipment, power, connectivity, and testing that ensures Disaster Recovery is achievable.

Disaster Recovery in Cloud


Disaster Recovery in Cloud entails storing critical data and applications in Cloud storage and failing over to a secondary site in case of any disaster. Cloud Computing services are provided on a pay-as-you-go basis and can be accessed from anywhere and at any point of time. Backup and Disaster Recovery in Cloud Computing can be automated, requiring minimum manual interventions.



3 Why use Cloud Disaster Recovery?

A Cloud Disaster Recovery provides numerous key benefits, over other types of disaster recovery strategies:

Easy Scale Up



- Departments can scale Cloud DR effortlessly, since it is very easy to increase the amount of resources that can be backed up in the Cloud by purchasing more cloud infrastructure capacity.
- Data centres on the other hand require sufficient server capacity to ensure high level of operational performance and allow data centre to scale up or scale out, depending on the Departments requirement.

Pay as you go



- With Cloud based DR, there is no need to invest upfront in hardware, or to pay for more infrastructure than the actual use at a given time. A Cloud-based Disaster Recovery service provides virtual machine snapshots of physical or virtual servers from the primary data centre. The Department can pay for storing the snapshots, application data in a suspended state, and replication of data from primary to the secondary (cloud DR) site for data synchronization. It has to pay for the infrastructure-as-a-service feature only in case of a disaster, wherein virtual machines (snapshots of primary servers) need to be brought online as a substitute for the primary site.
- While, a secondary physical DR site means investments in an additional data centre space, connectivity and servers, it leads to additional operational costs pertaining to power and cooling, site maintenance, and manpower requirements.

Geographic Redundancy



- Cloud-based Disaster Recovery makes it possible to leverage geographic redundancy features. This means that Departments can spread backed-up resources across multiple geographic regions in order to maximize their availability, even if part of the cloud that is used, fails.
- Whereas, on the other hand, it is costly to keep multiple DRs for same data in traditional DR setups

Faster Recovery



- With Cloud Disaster Recovery services, the DR site can be brought online within seconds or minutes—as opposed to a physical DR site. A virtual machine instance can be up and running within seconds. Typically, a physical DR site operates only during data replication, or in the event of an actual disaster. The time taken to make a DR site live will be more, in comparison to a Cloud DR. In addition, data loss is directly related to downtime. A Cloud DR site that boots up within a few seconds translates to data loss of just that timeframe.

4 Disaster Recovery Principles

<p>Distance</p>	<ul style="list-style-type: none"> • The distance for a DR site can vary depending on the types of disaster — such as earthquakes, floods, terror attacks, etc. The Departments should choose a DR location that fits its business model and regulatory requirements • Latency and performance of applications depends on distance in Disaster scenarios
<p>Recovery Time Objective (RTO)</p>	<ul style="list-style-type: none"> • RTO refers to the time an application can be down without causing significant damage to the business • Applications should be categorized by priority and potential business loss in order to focus on applications which are more critical first • Applications requiring near zero RTO require failover services
<p>Recovery Point Objective (RPO)</p>	<ul style="list-style-type: none"> • RPO refer to Departments data loss tolerance • Depending on application priority, individual RPOs typically range from 24 hours, to 12, to 8, to 4; down to near-zero measured in seconds • Near-zero RPOs will require continuous replication • 4-hour RPOs will need scheduled snapshot replication

5 Types of Disasters

A disaster can be related to any incident (both intentional and/or non-intentional) that causes severe damage to the operations and data of any Organization.

There are three major type of disasters:

Natural Disasters

(Hurricanes, earthquakes, floods, etc.)

Technological

(Chemical releases, power outages, etc.)

Man-Made

(Cybercrime, human error, terror attacks, etc.)

There can be various scenarios of disasters for which Departments should be prepared beforehand. The outages can range from a simple application failure to the disaster of whole Data centre. Below table shows some of the scenarios and he way Departments can deal with such outages.

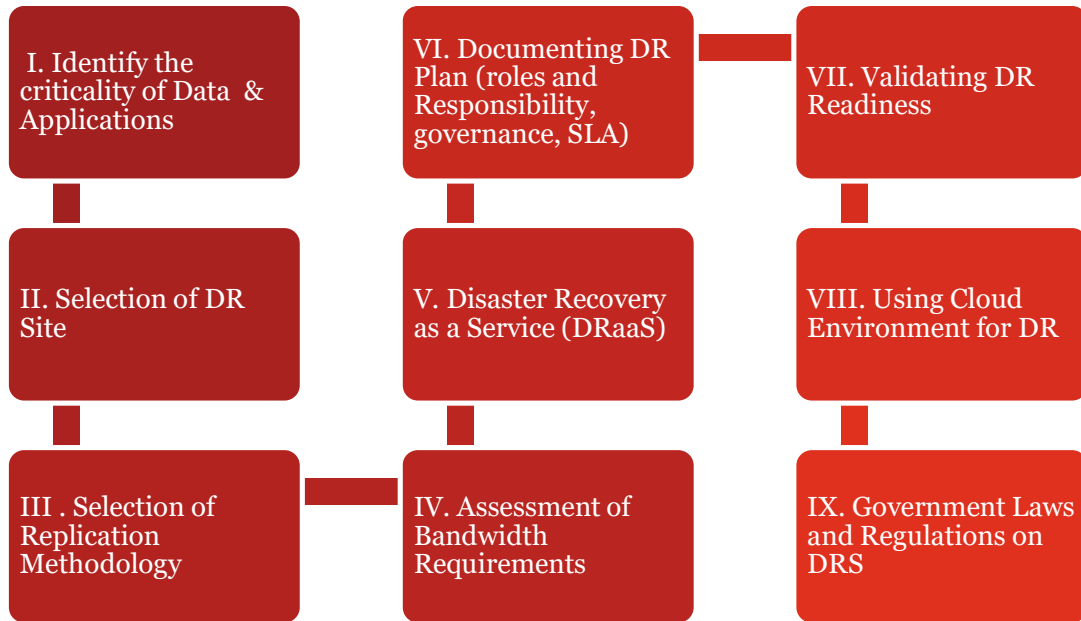
Application Failure	Network Failure	Data Centre Failure	Regional Failure
<ul style="list-style-type: none"> Application can fail due to changes in hardware/software configuration DR solution should have monitoring capabilities so that faults can be detected and alerts are sent to respective departments for action 	<ul style="list-style-type: none"> When users try to connect the on premise data centers through VPN, and for some reason network outage occurs In such cases, sufficient redundancy should be present for network connections. 	<ul style="list-style-type: none"> An unexpected event could affect an entire data center If multiple availability domains are present, it is recommended to deploy applications across the availability domains to accommodate potential issues for a particular data center If one availability domain is present, consider a combination of multiple fault domains and multiple-region configurations 	<ul style="list-style-type: none"> This is the most severe scenarios in cases of DR design It is recommended to use multiple cloud infrastructure regions. Backup or replication of data can be done in another region or a fully active-to-active standby can be setup in another region

Organizations can be categorized based on Disaster recovery planning:

No recovery plans	Such Departments fail to restore operations even during minor outages such a, power surge or server crash
Backup of data exist but there are no plans for Disaster management	In such cases, Departments need to back up their data regularly so that they can retrieve their data on the newly replaced systems in case of failure.
There is a backup data plan and external site to keep the backed-up data	Such Departments cannot tolerate to keep their systems down for an extended period. They have an arrangement to restore the required backed-up data which is kept at external site also called as data Off-siting
Remote, redundant sites as backup	Departments which have multiple data centres (at least two) that are located far away from each other. These data centres are interlinked with a strong communication network that facilitates the quick transfer of data in case of any disaster at either of these centres.
An exact replica of the working data system	This is where the data is backed up almost immediately per hour, per minute or even per second. With this method, Departments can recover from a disaster almost immediately. Even though this method is the most efficient, it is the most expensive as well.

6 Guidelines/Best Practices for adoption of Disaster Recovery

The Journey towards Disaster recovery setup provides step wise guidance in identifying the DR strategy suitable for their respective business. Below diagram depicts the steps which can be followed while planning for DR site.



Journey towards DR

6.1 Identify criticality of Data & Applications

Before implementing Disaster Recovery Site, it is important to classify, and group applications based on criticality. Such grouping of applications will help Departments to distinguish line of applications from each other in terms of their importance to the Departments, as well as their relative scope of influence on them.

MeitY shall launch IGCSF Toolkit, as a part of Risk & Security Assessment Decision making framework, which will help Departments to categorize their critical applications. The impact is divided into three categories, viz.

1. Assessment of impact on Departments (Tangible), in case of security breach
2. Assessment of impact on Departments (Intangible), in case of security breach
3. Assessment of impact on individual, in case of security breach

Based on the impacts (high, medium and low), Departments can categorize their respective application.

Categorizing business requirements based on priorities should be finalized. The below classifications will detail out the baseline for decision-making matrix.

Criticality Level	Failures of applications in this class can result in:
Mission Critical data/applications	<ul style="list-style-type: none"> • Widespread stoppage of applications with significant impact on Government operations • Public, wide-spread damage to Government reputation
Essential data/applications	<ul style="list-style-type: none"> • Direct impact on operations • Direct negative user satisfaction • Compliance violation • Non-public damage to Government reputation
Core data/applications	<ul style="list-style-type: none"> • Indirect impact on operations • Indirect negative user satisfaction • Significant Government department productivity degradation
Supporting data/application	<ul style="list-style-type: none"> • Moderate Government department productivity degradation

In addition to determining the criticality of applications, it is also necessary to understand the criticality of the Departments data

The Departments data remains equally critical as the data has evolved fast from mere excel or spreadsheet records to representing communication such as e-mail and important

Disaster Recovery Best Practices in Cloud Computing

digital documents. However, not all data in an enterprise is mission critical. It is important to classify data and define the associated metrics for retention, retrieval and archival. Missing this can increase costs exponentially (storage, backup, management, etc.). Classification helps in narrowing down the actual data that needs to be recovered in the case of a disaster.

Low Impact	All data and systems that does not require immediate restoration for the Departments to continue its operations
Moderate Impact	All data and systems that are important Departments can operate but in a diminished state
High Impact	All data and systems without which Departments operation can come to a halt.

User Departments should classify application and data based on criticality, as all the data and applications cannot be mission critical.

A survey that Forrester conducted in 2017 found that only 18% of organizations use either DRaaS (Disaster Recovery as a Service) offerings or public cloud IaaS offerings. On the other hand, Gartner estimates that the size of the DRaaS market will exceed that of the market for more traditional subscription-based DR services by 2018.

6.2 Selection of DR Site Architecture

Based on the criticality of applications and data, User Departments need to determine the best suited Disaster recovery site for their respective operations and perform evaluation of cost for selection of type of DR Site.

User Departments can select between internal and external Disaster recovery sites, based on their respective requirements. Below diagrams depicts the major difference between the two sites.

Internal Disaster Recovery Site:

When to use: Require aggressive RTO, require control over all aspects of the DR process.

Considerations: Expensive than an external site, Internal site needs to be built up completely by the Department

External Disaster Recovery Site:

When to use: When Departments require cost effective DR Sites.

Considerations: An outside provider owns and operates an external DR site.

3 types: Hot site, Warm site, Cold site

Distance is a key element in disaster recovery. A closer site is easier to manage, but it should be far enough that it's not impacted by the drive up and staff costs.

External Disaster Recovery Site

Hot Site

- Used for business-critical apps
- Fully functional DC
- Ready in the event of disaster
- It can be of 2 types:
 - Active-Active- Both sites are live
 - Active-Passive- Data is replicated in passive site

Warm Site

- Data is replicated but servers may not be ready
- Takes time to bring up servers to recover application in warm site
- Designed to be used for no- business critical apps

Cold Site

- Not ready for automatic failover
- High risk of data loss
- May take weeks to recover, as data from backup have to be loaded into servers
- Minimal infrastructure

6.3 Selection of Replication Technology

Data Replication is a way to ensure that Departments are prepared for disasters. Replication creates copies of data at varying frequencies, depending on the data in question and the industry of the organization backing it up. In the event of a disaster, the primary systems failover to this replicated system.

There are majorly two types of data replications:

1. Synchronous Replication

- Copies of data is created in real time on secondary site and locally
- Business continuity
- Very Low RTO RPO
- DC and DR sites in close proximity
- Minimize downtimes and assure a high infrastructural availability
- Limits:
 - The two sites cannot be far from each other

2. Asynchronous Replication

- It creates copies of data as per defined schedule
- It is suitable for Departments that can endure longer RTOs.
- DC and DR sites are distant apart*
- It allows to protect business even in case of large-scale disasters which may damage both sites (for instance, an earthquake)

*The distance between DC and DR is recommended to be minimum 100 kilometers.

Replication methodologies can also be controller based. Some of the methodologies in controller-based replication are mentioned in figure below:

	ARRAY BASED	APPLICATION BASED	HOST BASED	NETWORK BASED
Support of Heterogeneous Apps	Low; works for similar arrays	Low; works for specific application only	High; Storage agnostics	High; storage array & platform agnostics
Performance & Scalability	Good in high end arrays	Good for single Application platform	Workload is spread across servers	Good
Cost	High Entry Cost	High Entry Cost	Low entry cost; cost is based on no. of servers	High end entry cost; required intelligent switches
Replication Modes	Synchronous & Asynchronous	Synchronous & Asynchronous	Asynchronous	Synchronous & Asynchronous

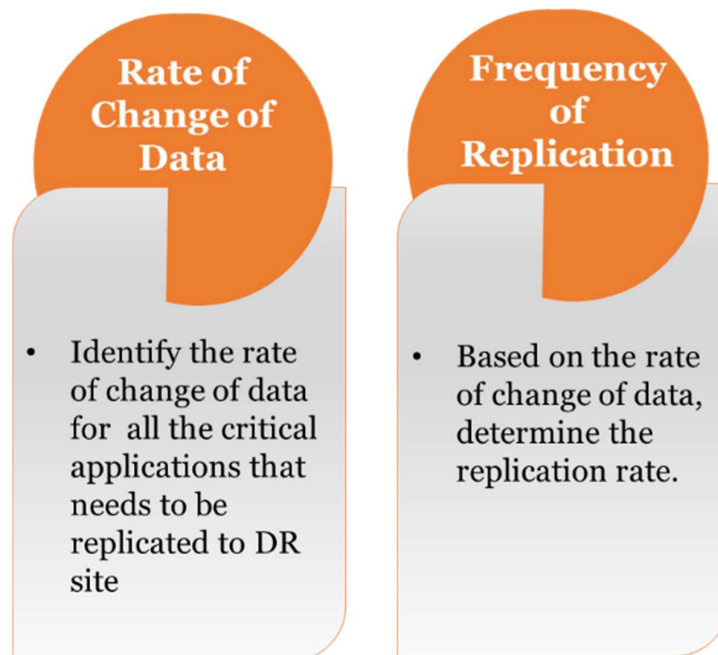
Departments that rely on mission critical data and cannot compromise on RTOs, can effectively leverage synchronous replication, while Organizations that can endure longer RTOs but need cost effective disaster recovery can use asynchronous replication. Also, Application based replication is the least preferred replication option due to dependency on individual Application vendor.



6.4 Understanding Bandwidth Requirements

Bandwidth and latency are equally critical as other factors while planning Disaster Recovery. Departments which replicate data for potential failover, both locally and remotely, should take bandwidth requirements into account while planning the DR site. The planning phase of a cloud-based DR implementation involves not only calculations with regard to keeping the off-site data up-to-date and within SLAs, but also with regard to user traffic when an actual recovery is needed. It is important to have data reside closer to its respective user departments as well as the applications or workloads which are being accessed.

There are two major factors which impact bandwidth requirement decision. Figure below explains the factors



The major considerations while estimating Bandwidth requirements while planning a DR site are:

- *While transferring data to the Cloud, sufficient bandwidth is required. Hence based on the application and data capacity and criticality, Departments need to specify the estimated bandwidth requirement.*
- *Department needs to specify the requirement of redundant network connectivity between DC and DR site.*
- *It is necessary to determine the network bandwidth requirements in Disaster scenarios, making the data accessible to its users after occurrence of a disaster.*

6.5 Disaster Recovery as a Service (DRaaS)

Since it is expensive to maintain a dedicated DR site, User Departments can choose to outsource this cost. Replacing the cost of dedicated site with a predictable expense is comparatively better option. Disaster recovery as a Service (DRaaS) enables full replication and backup of all cloud data and applications while serving as a secondary infrastructure. It actually becomes the new environment and allows an organization and users to continue with daily operations while the primary system undergoes restoration.

Reasons to consider Disaster Recovery as a Service (DRaaS) over On-premise DR Site:

1. **On-demand provisioning:** All cloud services offer on-demand self-service functionality. Once the service is initiated by the user, it takes only few minutes to get commissioned, which is much faster than commissioning the same service on premise.
2. **Easy Scalability:** Cloud services can be scaled exponentially. Adding resources to a cloud-based solution takes very less time and effort. On the other hand, if on premise DR is present, then Departments should be sure about the capacity in order to provide an adequate DR coverage.
3. **Removes maintenance overhead:** As Cloud platforms are part of managed service, maintaining and upgrading the underlying infrastructure is the responsibility of the Cloud service provider. If internal DR environment is present with the Departments, it also needs to be upgraded with the new features and patches, which require overhead of IT resource.
4. **DRaaS is cost effective:** Cloud Services works on pay per use model which is extremely cost effective as Departments can diligently control their spending by consuming and only paying for the resources they use. In case of on-premise environment, resources are often underutilized and do not run at full capacity. Enough hardware procurement is also an overhead cost.
5. **Multisite:** Resources can be replicated to many different sites to ensure continuous backup in the event of unavailability of one or more sites.
6. **Array agnostic:** DRaaS replicates any environment and is not service provider specific.

DRaaS pricing structure: DRaaS is often made up of several pricing components, including:

- Replicated data storage cost
- Software licensing costs (for disaster recovery and business continuity software to provide data replication)
- Computing infrastructure cost
- Bandwidth cost

Some DRaaS providers only charge for storage and software licensing when the service is not actually being used, adding compute infrastructure and bandwidth costs if the service is activated in the case of a disaster. Others charge for all components in the form of a "service availability fee," regardless of whether or not the service is actually used.

Disaster recovery Management Tool - Disaster recovery management tool is a part of DRaaS solution. It helps an Organization to maintain or quickly resume its mission-critical functions after a disaster. It is used to facilitate preventative planning and execution for catastrophic events that can significantly damage a computer, server, or network. It allows an organization to run instances of its applications in the provider's cloud. The obvious advantage is that the time to return the application to production, assuming networking issues can be worked out, is greatly reduced because there is no need to restore data across the Internet.

Some of the key features of a disaster recovery software are:

- Ease of use
- Monitoring capabilities
- Automatic backup of critical data and systems
- Quick disaster recovery with minimal user interaction.
- Flexible options for recovery
- Recovery point and recovery time objectives
- Compatibility with physical servers
- Easy billing structure
- Options for the backup target

While selecting DRaaS, Departments should consider the following:

DRaaS works on pay as you go model, so Organizations should select Service providers which provide different DRaaS service for different classes of applications.

In case of non-availability of Disaster recovery setup with primary CSP, services can be availed from other empaneled CSP's.

6.6 Documenting DR Plan

While documenting DR Plan, Departments should take a holistic view and focus on recovering the application services and not just servers. The technical recovery plan for each application/ service should be documented in a way that all the activities that need to be performed during recovery should be defined in a sequential manner.

- Design for end to end recovery
- Define recovery goals
- Make tasks specific: To make the system up and running, all steps should be pre-defined. Guess work should not be done. Documenting the steps is needed
- Maintain more than one DR recovery paths

It should cover all details such as physical and logical architecture, dependencies (inter- and intra-application), interface mapping, authentication, etc. Application dependency matrix, interface diagrams and application to physical/virtual server mapping play an important role in defining how applications interact with each other to deliver various functionalities.

6.6.1 Roles and Responsibility

Roles and responsibility should be clearly defined while planning for a Disaster Recovery Site. It should contain a governance structure often in the form of a Business Continuity-Committee that will ensure senior management commitments and define senior management roles and their respective responsibilities. The team composition should include below:

- **Disaster Recovery Planning (DRP) Coordinator:**

The DRP Coordinator shall have comprehensive decision-making powers, member from the higher Authority expected to lead the DR activities.

- **Crisis Management Team (CMT):**

The Crisis Management Team shall comprise of Management level personnel who shall analyze the damage at DC, advise the DRP Coordinator for Disaster Declaration, and initiate the recovery of Operations at the DR Site.

- **Damage Assessment Team (DAT):**

The Damage Assessment Team shall comprise of a management and technical expertise mixture of personnel who shall assess & report the damage at DC and take steps to minimize the extent of the same.

- **Operations Recovery Team (ORT):**

The Operations Recovery Team shall comprise of a management and technical expertise mixture of personnel who shall undertake the recovery operations for SDC at the designate DR Site.

The Business Continuity Committee will be responsible for:

Clarify their roles of all the members of the committee	✓
Oversee the creation of a list of appropriate committees, working groups and teams to develop and execute the plan	✓
Provide strategic direction and communicate essential messages	✓
Approve the results of Business Impact Analysis	✓
Review the critical services and products that have been identified	✓
Approve the continuity plans and arrangement	✓
Monitor quality assurance activities	✓
Resolve conflicting interests and priorities	✓

Roles and responsibilities of the Business continuity Committee should be clearly defined and well communicated in the Departments.

6.6.2 Segregation of Responsibility between CSP, MSP and a User Department

The segregation of roles and responsibilities between a Department, MSP and CSP can be seen in the below mentioned matrix:

Disaster Recovery Management	On-Premise	IaaS	PaaS	SaaS		
Program Management	MANAGED BY USER DEPARTMENT AND MSP					
Integration with Business Continuity						
Plan Maintenance						
Management Actions (Escalations, Declaration and Orchestration)						
Define application interdependencies						
Determine sequence of recovery						
Requirements definition (RTO, RPO)						
Application Validation					MANAGED BY MSP	
System Recovery						
Applications					MANAGED BY MSP	
DR Testing						
Database	MANAGED BY MSP					

Disaster Recovery Management	On-Premise	IaaS	PaaS	SaaS
------------------------------	------------	------	------	------

User Departments should focus on recovering the application services and not just servers. This is where sequence of the activities to be performed for restoring the operation plays an important role. The technical recovery plan for each application and service needs to document all the details of the activities that need to be performed along with the sequence of the activities.

6.6.3 Scope and Dependencies

Determining the most important VMs and including them into the recovery scope can help achieve shorter recovery time objectives. These VMs should be housing business-critical information, applications. Also, dependency links between these VMs, applications, and IT systems should be considered. For example, the operation of a particular application can be dependent on information housed on a different VM or vice versa. Dependencies also exist between employees and the components of the infrastructure. Figuring out and documenting such dependencies is necessary so that the Departments can continue their work with minimal interruptions.

6.6.4 Service Level Agreement

Service Level Agreement (SLA) as already detailed out in "Guidelines for User Departments on Service Level Agreement for procuring Cloud Services" on MeghRaj portal "https://meity.gov.in/writereaddata/files/Guidelines_User_Department_Procuring_Cloud%20Services_Ver1.0.pdf" can be referred to get more clarity on which SLAs to be negotiated while finalizing the Cloud service offerings. There are some key parameters which should be taken care of for managing Disaster situations well:

- Availability of the Service (uptime)
- Latency or the response time
- Service components reliability
- Each party accountability
- Warranties

6.6.5 Validating DR Readiness

The readiness assessment will help in evaluating the current status of Disaster recovery. Whenever downtime happens, for any reason, data/applications become unavailable to the users, which leads to an abrupt halt of all the functions.

The replicated resources to be **reviewed based on sector wise requirements (At least twice a year (six months))**, to ensure resources are effectively replicated and identified resources are in alignment with the business priorities and goals. The list below includes elements that should be reviewed on a regular basis to support critical server definition requirements.

- DR Drill
- Business impact analysis and risk assessment
- Application dependencies and interdependencies
- Backup procedure
 - Offsite storage for essential records
 - Data retention policies
- Recovery time objectives (RTO)
- Recovery point objectives (RPO)
- IT and Senior management signoff

Departments to plan for thorough testing the DR plan to unearth issues such as hardcoded IP addresses, host file entries, license file/ key, configuration details, dependency on other applications/ services, etc., resulting in the need to update the DR plan to make it robust. While it might not be possible to test all key business applications during a DR drill, it is important to focus on recovering the application(s)/service(s) and not just on servers. In essence, an untested DR plan is only a strategy.

6.6.6 Hybrid Cloud DR Scenarios

Departments can have Hybrid environment that use a mix of on-premises, private cloud and third-party, public cloud services with orchestration between the two platforms. But there can be differences between these platforms in term of infrastructure, hardware, software or configuration. So, having portability of the workloads into heterogeneous environments becomes a huge consideration. Besides, applications which are running on Cloud may have resiliency designed into them, but legacy applications and other assets might require working with service providers on various aspects like VPN access or frame relay connections in the event of a disaster.

Applications and environment can be rebuilt but if Departments loose data during a disaster it becomes difficult to return to normal operations. So, the foremost important thing is securing data. Efficient data replication technologies go a long way toward addressing that concern.

There are three major scenarios. These have been explained in detail in GI Cloud Architecture Guideline.




Figure 5

6.6.7 Government Laws and Regulations on Disaster recovery Site

The User Departments should also consider any guidelines which have been published by Regulatory authorities with perspective of Business Continuity or Disaster Recovery.

- While choosing the DRS and replication methodology, Department needs to follow the sector-wise guidelines
- DR setup/drills should be followed based on sector-wise regulatory requirements
- The manpower deployed at NS/ DRS should have same expertise as for Primary Data Centre in order to make DRS/NS functional in short notice, independently.

- **Failover Platform** - Maintaining an unused infrastructure in case of disaster is very expensive, so a better way is to have a public cloud provider in place which can act as a stand-by.
- **Dealing peak loads** – When departments face unpredictable events that lead to an application having higher workload than normal the hybrid cloud can come to rescue. It allows these application workloads to be provisioned at will.
- **Application's Test server**- Public cloud environments are ideal for testing. Test environment can be matched to the live ones closely. Hence, eliminating the need for a separate infrastructure for testing.

Annexure 1

Terminologies related to DR

Some of the key terms used for Disaster Recovery are described below:

- **Disaster Recovery** - The process of restoring and maintaining the data, systems, applications and other technical resources on which a business depends.
- **Application Recovery** - A component of Disaster Recovery that deals with the restoration of business system and data, after the operating system environment has been restored or replaced.
- **Disaster recovery as a Service** - Disaster recovery as a service (DRaaS) is the replication and hosting of physical or virtual servers to provide failover in the event of a man-made or natural catastrophe.
- **Business Continuity** - A system of planning, for recovering and maintaining both the IT and other functions within a Department regardless of the type of interruption. In addition to the IT infrastructure, it covers people, workplaces, facilities, equipment, processes, etc.
- **Business Impact Analysis** - A collection of information on a wide range of areas from recovery assumptions and critical business processes to interdependencies and critical staff that is then analyzed to assess impact a disaster may have.
- **High Availability** - High availability describes a system's ability to continue processing and functioning for a certain period of time - normally a very high percentage of time, for example 99.999%. High availability can be implemented into an Organization's IT infrastructure by reducing any single points of failure using redundant components.
- **Hot Sites versus Remote Sites** - A disaster recovery hot site is a remote physical location where you can maintain copies of all of your critical systems, such as data, documents, etc. A remote site provides a secondary instance or replica of IT environment, without office infrastructure. It can be securely accessed and used remotely by respective Organizations, through standard Internet connections from anywhere.
- **Mission-critical** - A computer system or application that is essential to the functioning of business and its processes.
- **Production or Primary Site** - The primary site contains the original data that cannot be recreated.
- **Recovery Time Objective (RTO)** - The RTO is the duration of time and service level within which a business process must be restored after a disruption in order to avoid unacceptable losses. RTO begins when a disaster hits and does not end until all systems are up and running. Meeting tighter RTO, windows requires positioning secondary data so that it can be accessed faster.
- **Recovery Point Objective (RPO)** - The RPO is the point in time to which an Organization must recover data. The RPO is the "acceptable loss" determined by an

Disaster Recovery Best Practices in Cloud Computing

Organization.in a disaster situation. The RPO dictates which replication method will be required (i.e. backups, snapshots, continuous replication).

- **Redundancy** - A system of using multiple sources, devices or connections so that no single point of failure will completely stop the flow of information.
- **Risk Assessment** - The identification and prioritization of potential business risk and disruptions based on severity and likelihood of occurrence.
- **Secondary Site (or DR Site)** - The secondary site contains information and applications that are built from the primary repository information. This site is activated whenever the primary site becomes unavailable.
- **DR Drill:** DR Drill is a routine activity done by an organization to check if there is business continuity in case if the DC site is down due to an unexpected event.

Cost considerations for setup of Disaster Recovery*

Available Options	Type of Disaster Recovery Setup	Cost	Impact
Option - 1	<ul style="list-style-type: none"> • Site would be Hot or warm Disaster recovery Site • Disaster Recovery resources are configured at 100% of Production Data Centre capacity. • All Application, Database, Stateless and IT infrastructure servers are replicated to DR site and are operational at all the time. • Recovery point operation is High/Aggressive (near zero to Minimal data loss) • Recovery time operations may be defined based on department needs (2 to 24 Hours for operations availability) 	High Cost Model	<p>Pros:</p> <ul style="list-style-type: none"> • Departments can achieve aggressive RPO and RTO's • Good performance at DR Site • Automatic failover can be achieved • Minimum intervention to operationalize the DR Site. <p>Cons:</p> <ul style="list-style-type: none"> • High Cost and requirements of physical site/It Infrastructure/Cloud resources equivalent to Production DC site • High Bandwidth and Data Transfer cost
Option - 2	<ul style="list-style-type: none"> • Site can be hot or warm Disaster recovery Site • Disaster Recovery resources are configured at 80% to 50% of Production Data Centre capacity. • Selective application, Database, Stateless and IT infrastructure servers are replicated to DR site and functional at the time. 	Moderate Cost Model	<p>Pros:</p> <ul style="list-style-type: none"> • Data can be replicated based on criticality and department needs • Cost effective model, as only required IT Infrastructure/Cloud resources needs to be procured <p>Cons:</p> <ul style="list-style-type: none"> • High time for RPO and RTO

Disaster Recovery Best Practices in Cloud Computing

Available Options	Type of Disaster Recovery Setup	Cost	Impact
	<ul style="list-style-type: none"> IT Infrastructure resources may be downsized at Disaster recovery site as compared to Production Data centre Recovery point operation is minimal to moderate (2 to 24 Hours of data loss) Recovery time operation may be defined based on department needs (4 to 48 Hours for operations availability) 		<ul style="list-style-type: none"> Downtimes may be required to operationalize the complete system at Disaster Recovery Site Performance degradation may be faced due to availability of lesser resources Require high time to procure and installation of resource to acquire required performance
Option - 3	<ul style="list-style-type: none"> Site can be hot or warm or Cold Disaster recovery Site Disaster Recovery resources are configured with less than 50% of Production Data Centre capacity. Minimum DB servers and application services are replicated to DR site. IT Infrastructure resources may be downsized at Disaster recovery site as compared to Production Data centre Recovery point operation may be minimal to moderate (2 to 24 Hours of data loss) Recovery time operation may be defined based on department needs (4 to 48 Hours for operations availability) 	Low Cost Model	<p>Pros:</p> <ul style="list-style-type: none"> Low cost as minimum IT Infrastructure/Cloud resources are procured, configured and replicated to DR site <p>Cons:</p> <ul style="list-style-type: none"> High time for RPO and RTO Downtimes may be required to operationalize the complete system at Disaster Recovery Site Performance degradation may be faced due to availability of lesser resources Require high time to procure and installation of resource to acquire required performance
Option-4	<ul style="list-style-type: none"> Site would be Cold Disaster recovery Site Only Data is replicated to Disaster recovery site Only Storage capacity is configured at Disaster recovery site RPO can achieved as per requirement, but RTO would be very high 	Minimal Cost model	<p>Pros:</p> <ul style="list-style-type: none"> Only Storage cost is required <p>Cons:</p> <ul style="list-style-type: none"> Disaster Recovery site would be Cold Site and would be non-operational, due to unviability of resources Require high time to procure and installation of resource to bring up the application

**Note: provided tables content is listed based on best practices, configuration and setup may vary as per department needs*