MAXIMISE UPTIME, MINIMISE DATA LOSS

While cloud storage has been proven more cost-effective and more reliable than on-premise approaches, businesses still need to make-use the features built-into platforms like Azure so they're able to benefit no interruption in service and maximum mitigation against data loss.

The great news is putting a disaster recovery plan in place for cloud isn't nearly as arduous as what it is in the on-premise world.

Follow this guide and you'll be well on your way to understanding the principles of mitigating against disaster in the Azure environment.

STEP ONE: CHOOSE AN APPROPRIATE REDUNDANCY OPTION

Azure Storage by default maintains multiple copies of your storage account to ensure durability and high availability. Which redundancy option you choose for your account depends on the degree of resiliency you need for your applications and of course, budget.

The options to choose from comprise:

- Locally Redundant Storage (LRS)
- Zone Redundant Storage (ZRS)
- Globally Redundant Storage (GRS)
- Globally Zone Redundant Storage (GZRS)

With LRS three copies of your storage account are automatically stored and replicated within a single datacentre. This means, should there be a hardware failure, there's two redundant copies of your data within the same datacentre to failover to.

With ZRS a copy of your data is stored and replicated in three separate availability zones within the same region. So should there be a catastrophic failure datacentre wide, you're able to failover to another datacentre in the same region.

Recovery of a single copy of a storage account occurs automatically with LRS and ZRS.

With globally redundant storage (GRS, GZRS, and RA-GZRS), Azure copies your data asynchronously to a secondary geographic region at least hundreds of kilometres away.

This allows you to recover your data if there's an outage that spans the primary region.

A feature that distinguishes globally redundant storage from LRS and ZRS is the ability to fail over to the secondary region if there's an outage in the primary region. When this happens, the region that's failed over to becomes the new primary endpoints for your storage account.

RA-GRS and RA-GZRS redundancy configurations provide geo-redundant storage with the added benefit of read access to the secondary endpoint if there is an outage in the primary region.

If an outage occurs in the primary endpoint, applications configured for read access to the secondary region and designed for high availability can continue to read from the secondary endpoint.

Microsoft recommends RA-GZRS for maximum availability and durability of your storage accounts.

You can find our more about this in detail here.

STEP TWO:
PLAN FOR THE EVENTUALITY OF A FAILOVER

When it comes to failover, there's really only two concepts to grasp, namely customer-managed failover and Microsoft-managed failover.

The names of these go a long way towards explaining their purpose, suffice it to say that with the former, customers can manage storage account failover if there's an unexpected service outage, whereas Microsoft-managed failover will take place only in a case where a severe disaster occurs in the primary region.

Best practise dictates that a disaster recovery plan should be based on customer-managed failover, while Microsoft-managed failover will be initiated automatically in extreme circumstances.

With customer-managed failover, If the data endpoints for the storage services in your storage account become unavailable in the primary region, you can fail over to the secondary region. After the failover is complete, the secondary region becomes the new primary and users can proceed to access data in the new primary region.

To fully understand the impact that customer-managed account failover would have on your users and applications, you can read more about the failover and failback process **here**.

In extreme circumstances where the original primary region is deemed unrecoverable within a reasonable amount of time due to a major disaster, Microsoft may initiate a regional failover.

In this case, no action on your part is required. Until the Microsoft-managed failover has completed, you won't have write access to your storage account. Your applications can read from the secondary region however, if your storage account is configured for RA-GRS or RA-GZRS.

STEP THREE:
EXPECT DATA-LOSS AND INCONSISTENCIES

Even in the cloud era, some data loss should be expected, especially in environments where loads of write activity takes place. This is because the data is written to the secondary regions asynchronously. There's always a delay before a write to the primary region is copied to the secondary.

If the primary region becomes unavailable, the most recent writes may not yet have been copied to the secondary.

When a failover occurs, all data in the primary region is lost as the secondary region becomes the new primary. All data already copied to the secondary is maintained when the failover happens. However, any data written to the primary that hasn't also been copied to the secondary region is lost permanently.

This is why it's important to take note of the Last Sync Property.

As its name suggests, the Last Sync Time property indicates the most recent time that data from the primary region is guaranteed to have been written to the secondary region.

All data and metadata written prior to the last sync time is available on the secondary, while data and metadata written after the last sync time may not have been written to the secondary and may be lost. Use this property if there's an outage to estimate the amount of data loss you may incur by initiating an account failover.

There's a more detailed explanation of some more specific considerations **here**.

STEP FOUR:
CALCULATING THE TIME AND COST OF FAILING OVER

The time it takes for failover to complete after being initiated can vary, although it typically takes less than one hour.

It's important to note that a customer-managed failover loses its geo-redundancy after a failover (and failback). Your storage account is automatically converted to locally redundant storage (LRS) in the new primary region during a failover, and the storage account in the original primary region is deleted.

When you re-enable geo-redundant storage (GRS) or read-access geo-redundant storage (RA-GRS) for the account, be aware that converting from LRS to GRS or RA-GRS incurs an additional cost.

The cost is due to the network costs of re-replicating the data to the new secondary region.

The re-replication time depends on many factors, such as the number and size of the objects in the storage account, the CPU, memory and disk resources available for background replication, if the storage account contains blobs, the number of snapshots per blob and whether or not the storage account contains.

STEP FIVE:
DESIGN FOR HIGH-AVAILABILITY

It's important to design your application for high availability from the start. Refer to these Azure resources for guidance in designing your application and planning for disaster recovery:

- **Designing resilient applications for Azure**
- **Resiliency checklist**
- **Using geo-redundancy to design highly available applications**
- **Building a highly available application with Blob storage**