silverbug

Disaster Recovery: The FAQs

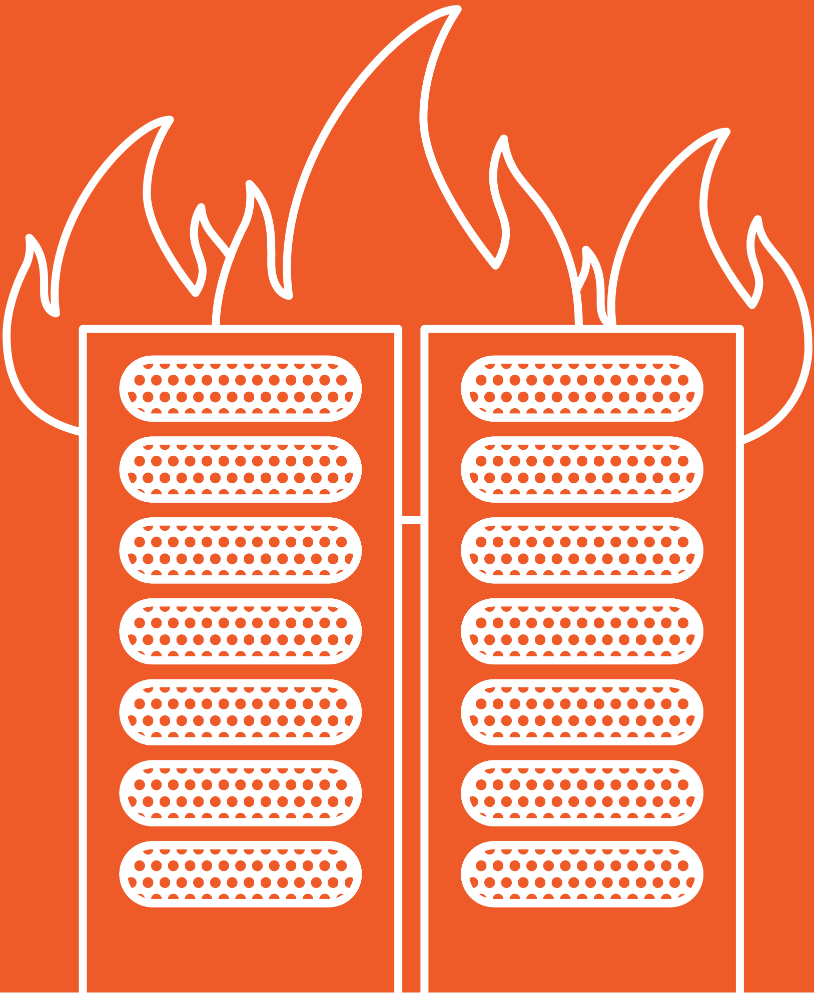

Everything You Need to Recover Your Data If Disaster Strikes

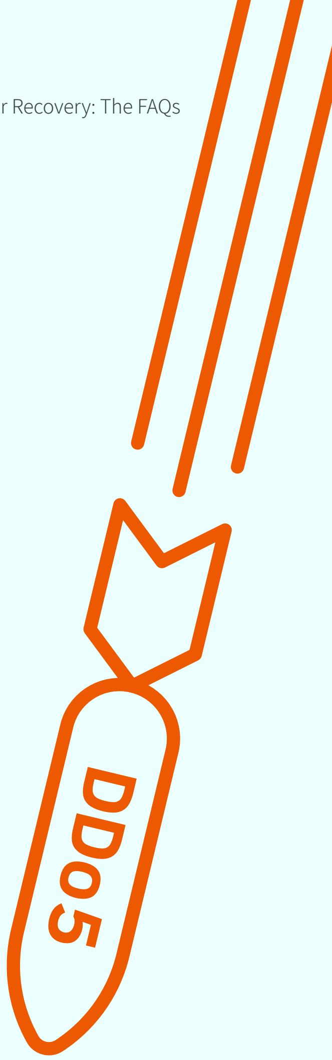

Introduction

Disasters can strike in many different forms, from physical accidents like a burst water pipe to Ransomware that encrypts all your vital data.

There are various options for Disaster Recovery, with cloud-based methods being one of the most popular. Cloud-based recovery plans are known for being convenient, flexible and easily maintained, among many other benefits. But many people do question their security. How secure can data really be when it's stored on the internet via a cloud server? And of course, there are some who wonder if Disaster Recovery Plans are even worth having.

This whitepaper will list and answer some of the most frequently asked questions about Disaster Recovery, from what it is and its background to the importance of backups and the security of cloud-based data.

We'll explore the ins-and-outs of each in more depth too.

1. Why do I need Disaster Recovery?

Because it will help you restore as much data as possible in the event that you experience partial or total data loss.

2. What is the difference between backups and data replication?

A backup is a version of your data that has been replicated up to a particular point in time, while replication is an alternative copy of your data with all of the changes made to the original made in real time.

3. What is the difference between RPO and RTO?

RPO allows businesses to know up to what point in time their recovery plan can still go smoothly while RTO allows you to calculate how quickly your company needs to recover from data loss.

4. Aren't backups and archives the same thing?

No, backups are historical copies of your data that are often required instantly in case of data loss. Archives are copies of your data that need to be stored for decades.

5. How is Disaster Recovery in the cloud different to other methods?

Instead of being stored on a physical hard drive, they'll be stored on a virtual drive on the internet which you can access from anywhere.

6. I have a lot of data – won't the recovery time be too long?

Perhaps, but the amount of time it takes to recover data depends on the problem as well as the number of files.

7. How can I be sure the recovery plan will work when I need it the most?

You need to test your plan and ensure every employee is following the right steps to protect your data.

8. How do I know my data is secure?

Unfortunately, you'll never be 100 percent sure. But, you can take preventative measures to protect your data as much as possible.

9. Do I have enough bandwidth to support cloud backup?

Find this out by reviewing how many users are on your network and how much internet activity goes on. Then, match that with your bandwidth needs.

10. Where is my data actually stored?

Virtually, but the data still needs to be on a hard drive somewhere - usually on a server farm whose servers operate literally all the time.

11. How do I access my DR services?

You can use a VPN client to securely access your Disaster Recovery systems from any computer or use a secure shell (SSH) tunnel to connect your office to the Disaster Recovery systems.

12. What happens if I don't have DR?

If you don't have a plan or backups, then you could face huge, irreversible data loss.

13. How much does it cost?

Prices vary depending on the cloud-based service and the business needs, but cloud plans are more cost-effective than traditional services - mainly because you don't need to maintain the physical hard drives.

14. Who manages and monitors the DR service and replication?

You can do it yourself or set up automated Disaster Recovery monitoring tools to do it for you.

15. What are the WAN/Internet requirements for your platform?

You can access your Silverbug-hosted Disaster Recovery systems from any reasonable internet connection.

1.

Why do I need Disaster Recovery?

Because it will help you restore as much data as possible in the event that your business experiences partial or total critical data loss.

Let's break this down.

So, imagine walking into your office, bright and early on a Monday morning, only to be greeted with a "Your Data is Corrupt" message when you try and turn on your computer. Perhaps it's because you've been a victim of a ransomware attack and the cybercriminals have encrypted all of your data, including vital business information. Or maybe your entire office has been flooded and all of your hard drives have been destroyed and your backups rendered useless.

The point is, you've lost your data and your business can't function well or at all without it. This is where Disaster Recovery comes swooping in to save the day. If you didn't have this, where could you possibly begin to recover your critical data?

2.

What is the difference between backups and data replication?

These two terms are often used interchangeably and while there are some similarities, the main difference is how backups and data replications are made.

A backup involves making a copy or copies of data, which then require a place to store these 'tapes' of data - such as a flash drive or a cloud-based solution. The whole concept of backups are based on snapshot technology, which is where copies of data are taken at a predetermined point in time and stored in a particular order. Old backups may be replaced by new backups, but usually, each set will be treated as a separate piece of data.

Backups ensure that there is zero loss of critical data or production data in the event of a disaster, such as the data becoming corrupt or deleted. When this happens, a historical version will be needed in order to access a usable copy of the data. This is where backups come in - they ensure there is an intact, usable copy.

Data replication, on the other hand, is when data has been copied and transferred to another platform or drive. Unlike backups, which are historical versions of data, replicated copies are literally exact, word-for-word, code-for-code copies of the original file. Any changes made to your original copy will also be made to the replicated version in real time. That means if something goes wrong and all of your original data is erased, so will your replicated version. You can't rely on it as a backup.

Replication drastically reduces the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) of a company because it allows instant access to the exact copy of the data file you might have lost. This then means businesses experience very little data loss - as long as that replicated copy isn't corrupted.

However, replication requires both time and money. You need another identical system to your original platform to hold your replicated files. Plus, replication is time-consuming. This can then double your IT costs. In contrast, backups are a largely automated process so in most cases, you don't have to manually create them (unlike replicated files) - your files will be backed up automatically as you work.

3.

What is the difference between RPO and RTO?

RPO and RTO are two of the most important measurable factors in a Disaster Recovery Plan. They guide companies on which backup plan they should choose for optimal recovery.

Recovery Point Objective (RPO)

Recovery Point Objective (RPO) refers to the time between data backups and the amount of data that could potentially be lost in between these backups. This depends on the individual company and their capacity. For example, if they can afford to lose a day's work, RPO will often be set at 24 hours.

Basically, RPO allows businesses to know up to what point in time their Disaster Recovery Plan can still proceed smoothly, given the volume of data lost during that period of time.

Recovery Time Objective (RTO)

Recovery Time Objective (RTO) refers to the duration of time within which a company or business process must be restored after a disaster in order to avoid undesirable consequences (that come after a break in business continuity). The goal here is to calculate how quickly your company needs to recover. This can then help you determine what you need to prepare and how much budget you need to assign to restoring business continuity.

For example, if your RTO is five hours, that means your company can survive for this long with systems being down. Because it's a short period of time, you'll likely have to invest more time and money to ensure your systems can be recovered quickly. If your RTO is a month, you won't need to invest as much time and money.

SPEAK TO YOUR OWN DISASTER RECOVERY EXPERT

4.

Aren't backups and archives the same thing?

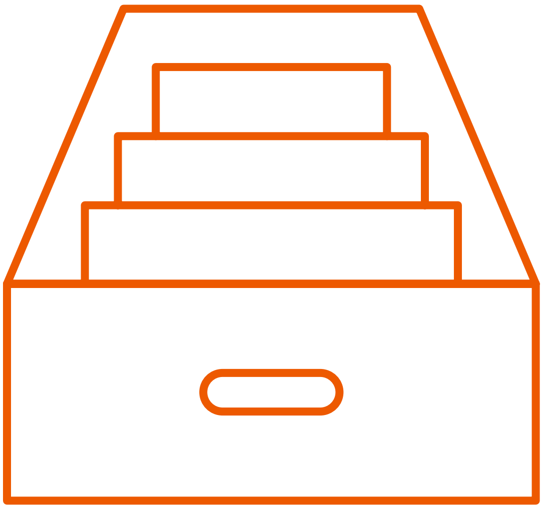

Like with backups and data replication, backups and archives are often thought to be similar because they both involve storing copied data. But the two were designed with different purposes in mind.

Like we mentioned earlier, backups are historical copies of data that are meant to act as a failsafe in the event of equipment failure, a virus attack or other catastrophe. Backups are often required instantly in emergencies so businesses can get back up and running as soon as possible.

Data archives, on the other hand, are meant to act as a repository for data that needs to be stored for long periods of time - such as decades. Unlike backups, they don't need to be accessed urgently because they're designed for long-term data retention.

5.

How is Disaster Recovery in the cloud different to other methods?

Cloud-based Disaster Recovery is different from the traditional method of transferring to a physical backup disk or drive. Instead of your backups being stored in a physical hard drive, they'll be stored in the cloud on the internet, so it's essentially a virtual hard drive.

One of the main benefits of a cloud-based plan, among others, is that it doesn't take up any physical space. This can make a cloud-based solution a much more cost-effective option because you don't have to buy more physical space to accommodate the extra storage systems you might need or spend money on electrical and maintenance bills. With a cloud-based solution, you just buy more virtual cloud space when you need more. That's all.

Another major benefit is that you can access these backup files wherever you are, thanks to the cloud location. Employees don't need to be physically present in your office - they can access your server from home or any other remote location, as long as they have access to WiFi. There are, of course, several precautions you should take if you want your cloud server to be as secure as possible.

6.

I have a lot of data – won’t the recovery time be too long?

The amount of time it takes to recover data depends on the problem, not just the amount of files. Generally, recovery should only take two to five days, but a Disaster Recovery specialist needs to take a look and diagnose the problem in order to get a better estimate of the time needed.

They will consider:

- The size of your hard drive - of course, the bigger the drive, whether physically or logically, the longer the recovery process as more time will be needed to correctly clone a larger device’s data.
- The model or series of the hard drive.
- The type and size of the files.
- The environment your computers were operated in - for example, hot environmental conditions cause more physical damage to drives, which can lengthen the recovery time.

7.

How can I be sure it will work when I need it the most?

You can be sure your Disaster Recovery plan will work by ensuring you follow some best practices. These include:

- Testing your plan. Like fire drills, you need to check that your Disaster Recovery plan will work by testing it and ensuring everyone is following the steps properly.
- Ensure that your Disaster Recovery tools work across a variety of platforms, applications and software. There's no point having a recovery plan if it's not compatible with your systems.
- Update the tools regularly. This ensures you always have the latest version of your recovery software so you're less likely to experience hiccups in your recovery process.
- Backup at regular intervals. In fact, backup all the time. You never know when disaster might strike and that one time you don't backup because "it'll be fine!" might be the one time a virus attacks and you lose all of your work. And without backups, you can't recover very well, if at all.
- Manage who has access to your sensitive files and their relevant backups. Only allow access to authorised members and keep a note of who can modify them.

SPEAK TO YOUR OWN DISASTER RECOVERY EXPERT

8.

How do I know my data is secure?

Truthfully, the short answer is that you don't. But you can take preventative measures to ensure your data is as well-protected as possible.

- Avoid sharing sensitive information in the cloud. It's like when you park your car - would you leave it with valuables on show?
- Ensure you know how your cloud-based storage system works.
- Choose secure passwords. Did you know that 90 percent of all passwords can be cracked by experts in seconds? Don't create passwords that are easy to guess and ensure only authorised members of staff have access to them.
- Encrypt your data. This is one of the best ways to protect your data.

9.

Do I have enough bandwidth to support cloud backup?

You need to find out what your network capacity is by looking at how many users are on your network and how much internet-required activity goes on - for example, emails sent and received, media streaming, uploads and downloads, Google searches and so on. Then, match that to your bandwidth needs. After all, a bigger bandwidth means less time waiting for downloads and uploads and more time being productive, doing non recovery-plan-related tasks.

You should also review your RPO, how often you want to send data offsite and how much of your bandwidth you're willing to allocate to this task.

10.

Where is my data actually stored?

In a nutshell, cloud computing involves taking data from your own personal little space and placing it in a special section of the internet. Nothing will be stored in your local hard drive, but you can still access this data from any location using any device at any time.

Although this data is stored virtually for you, it still needs to be physically stored on a hard drive somewhere. Companies that offer cloud-based services have huge server farms, which are essentially enormous, cavernous warehouses filled with servers which are running 24 hours a day, seven days a week and 365 days a year. Basically, all the time.

11.

How do I access my DR services?

You have a couple of options here. First, you can use a VPN client to access your Disaster Recovery systems from any computer. This will securely log you into your recovery systems as though you were in your office. This is a great option if you need to access your systems but you're not in office.

The alternative is to use a secure shell (SSH) tunnel to connect your office to the Disaster Recovery systems. Tunnelling is a protocol that allows for the secure movement of data from one network to another. In data security and recovery, it will allow easy access to your data storage centres. You should also have multiple backup tunnels which can become active whenever a primary tunnel fails.

12.

What happens if I don't have DR?

Well, unfortunately you'd be in a bit of pickle. Hopefully you should have at least some backups in place (if not, then you're in even more of a pickle). The best thing you can do is to consult a recovery specialist and ask them to work their magic. But there's no guarantee that this will work and it can be a costly service.

The lesson from this? Every business, big and small, needs to have a Disaster Recovery plan - you never know when you might experience a catastrophe and you need to think about how you'll recover from it, from data recovery to replacing damaged hard drives. You need to have a plan for business continuity.

13.

How much does it cost?

The price varies from cloud-based service to service and the business individual needs, but it's important to realise that cloud recovery plans are a much more cost-effective option than traditional services.

Many of the traditional Disaster Recovery services come at very high costs, particularly due to the costs of maintaining the backup hard drives which can accumulate over time. In contrast, cloud-based recovery allows businesses to store their backups in an offsite server, managed by an external company.

This means the price per server and price per square metre that the server takes up aren't factored in as prominently. So, you don't need to:

- Maintain the servers, which reduces hardware costs.
- Buy additional space. When you require more storage space, you just buy cloud space by upgrading your account or storage limits. This will be much cheaper than buying physical space.
- Invest in the utilities bills required to maintain the appropriate environmental conditions.
- Have operations staff to manage the servers.

Overall, a cloud-based recovery service is a much more cost-effective solution than the traditional methods.

14.

Who manages and monitors the DR service and replication?

You can do it yourself or you can set up automated Disaster Recovery monitoring tools that do this for you. These tools are extremely helpful because they can monitor the changes that happen in data storage environments and notify you of them so you can be ready for a potential disaster.

You could have:

- Tools that store information you've collected about your plan, which you can use to create planning documents.
- Tools that help you set up scenarios so you're prepared for a variety of circumstances. They can automatically create backups of your data for you.
- Tools that monitor the data processes that are going on - for example, creating ongoing reports on the process of data backups so you're constantly aware of what stage the backup process is at.

15.

How do I know my data is secure?

Because cloud-based services store data in a virtual 'cloud', they require an internet connection to be able to transfer this data. Of course, the stronger the connectivity and the higher the bandwidth, the faster your data can be backed up, uploaded onto the cloud and downloaded in the case of a disaster.

With a Silverbug-hosted Disaster Recovery system, you can access it from anywhere as long as you have a reasonable internet connection. This can work even if you were at home on a work computer.

Now, Speak to an Expert

If you haven't already done so, it's time to speak to a Disaster Recovery expert. Not doing so runs the risk of your business becoming just another company who never recovered from an IT disaster.

SPEAK TO YOUR OWN DISASTER RECOVERY EXPERT

silverbug

0845 304 8888 enquiries@silverbug.it
Offices in the UK (Milton Keynes, Leeds and London)
and in the US (Houston, Texas)