www.tsg.com

THINKING ABOUT THE UNTHINKABLE

A guide to assessing risk and creating a disaster recovery plan

TSG

For most businesses, it's unlikely that 'disaster' will result in loss of life, however, the consequences will inevitably include significant disruption and the financial implications could be crippling if suitable plans aren't in place to minimise the impact.

That's aside from the frustration, stress and mental anguish that goes with such challenging circumstances.

As expectation around 'on-demand' and 24/7 accessibility has grown over the last decade, businesses have become increasingly reliant on technology and in the broader context of business continuity, disaster recovery relates specifically to the data and IT systems that underpin business operation.

Learning lessons from those who are tackling disasters that wipe out entire communities, the disaster recovery cycle has four distinct yet interrelated phases: mitigation, preparedness, response and recovery.

These should all be included in any organisation's disaster recovery planning to define a well-structured method of restoring data and systems that for whatever reason have become inaccessible.

Much like car insurance, the disaster recovery plan is something that you hope will never be implemented but the economics are simple – in fact, if you apply good business practice, not only will your investment in disaster recovery planning pay back in terms of getting back on track, it could also minimise the risks in the first place.

Here are a number of areas that you might want to consider when you're thinking about the unthinkable.

1. Probability First

It sounds a bit obvious but...work out the most likely disaster events.

Sometimes it's easy to think about events such as the headline grabbing floods as a major risk, but if you aren't on a flood plain, near a major water course or your server equipment is all housed on the second floor, is flooding really a contingency you need to plan for?

However, a power outage may be far more likely – and in some cases might be caused by a flood.

Or you may be affected by a major telecoms fault like the one we saw back in 2004. Let's face it, even minor telecoms faults can cause chaos.

This particular incident was undoubtedly a freak event but weeks of disruption and downtime were the result when a serious fire in a cable tunnel in Manchester destroyed large sections of BT's fibre network.

Reports at the time suggested that businesses as far afield as Sweden were affected and one estate agency business counted losses into tens of thousands of pounds.

2. Like for Like

Group together similar types of event.

Once you know the types of disaster events you are planning for, try to put them together.

For example, the recovery from a sustained power outage and a local telecoms failure may invoke similar recovery plans.

If you're running services in the Cloud, then your contingency plan might be to send everyone home to work using remote access.

Grouping recovery plans together helps to reduce the costs of planning around many different disaster scenarios.

3. Not All Data is Equal

It's often the view that all services are absolutely critical and must be instantly recovered with zero data loss.

Realistically, it's unlikely that most businesses could afford to take this view. So, by considering the impact to the business of each individual service failure you can plan a far more cost effective – and economically viable – disaster recovery plan.

Every business is unique in its requirements and the key to this is to understand the RPO (Recovery Point Objective) and the RTO (Recovery Time Objective) for each service.

For each set of data, system or application work out two simple things:

1) How much data can you afford to lose and this is your Recovery Point Objective. It might be 5 minutes, 1 hour, 4 hours, a couple of days or none at all.

2) Now consider how quickly you need to get that service back online and this is your Recovery Time Objective. Again it could be 5 minutes, 1 day, 5 days or as close to instant as is possible

Work out what will have the biggest impact?

If you rely on email to receive the majority of your orders then email is arguably most critical.

If delivering great service relies on your contact management system, then you need both the application and the data available to the frontline team.

If losing your manufacturing application or data brings production to a grinding halt then that's clearly going to be a high priority.

It's the combination of RTO and RPO that will allow you to build a cost-effective and successful strategy for recovery.

4. Money vs. Reputation

It essential to understand both the financial and the less tangible impacts of a service failure.

Some impacts are obvious.

If you rely on your website to take £2,000 worth of orders per hour and it fail the calculation is simple. Every hour the site is down results in a £2,000 loss.

However, some are far less obvious.

Will the customer come back to your website and try again? Have you lost a potential regular customer and anyone they might have recommended or referred?

What about the impact on staff morale if you are losing orders or having to work late to catch up on work following major downtime?

It's clear that many businesses do not get it right when considering their disaster recovery plans, whether the cost is financial or to their reputation.

There are some pretty scary – and probably exaggerated - statistics around suggesting businesses that experience a disaster are considerably more likely to fail within a couple of years of the event. Unsurprisingly, the legendary figure of 80% is widely disputed!

5. Prioritise

Once you've identified the potential risks and the impact on different areas of the business, collate the information into a coherent format for review.

The next step is to prioritise based on their likelihood and business impact – and of course, the cost to mitigate against each risk profile.

A simple ROI formula plan will help and this is where help from your IT provider comes in. You can discuss options and budgets around each of the plans and formulate the final results into a coherent Disaster Recovery plan.

Clearly it must make economic sense. It's pointless investing £100,000 in a complex solution if the impact of the risk is only ever going to be £15,000.

6. Allocate Responsibility

Who makes the call?

One area that's often overlooked is allocating whose responsibility it is to call it a disaster in progress and invoke the DR plan.

This may sound obvious but the decision may involve multiple people or third parties. For example, is the server really totally irrecoverable or would a quicker option be to uplift and repair the systems rather than start workplace recovery procedures – which might involve much more complexity and cost.

Is it the local IT Manager who makes that call, the business owner, or the IT company? It's essential that everyone

understands their role and the role of others in implementing the DR plan.

And what if key people aren't available – who makes the call then?

7. Document it

Again, who's responsible for documenting the DR plan – someone within the business or your IT service provider?

Obviously, the document should be accessible in the event of a disaster – so copies should be held in multiple locations – and more than one person in the organisation should be aware of where to locate the plan and how to invoke it.

The DR plan should be broken down in to different types of disaster that need to be planned for, with the relevant recovery steps, and the contact details of everyone involved in implementation including the decision makers.

Not only should the plan be cost effective but it should also give the peace of mind that, should the unthinkable happen, you have considered all eventualities. Above all, it must be clear and simple to implement, especially given the levels of stress that will already have been created by the situation that's led to it being invoked.

8. Failback

Clearly the main element of any disaster recovery plan is the failover procedure – i.e. how you go from a disaster situation to having usable services again.

However, it's also important to consider the failback to normal systems, especially as this could have a significant financial impact.

Part of the broader business continuity plan – i.e. covering general operational issues rather than just the data and systems recovery element – may be to send everyone to a managed office with remote access to their systems.

But how long are you prepared to do this? What will it cost?

And how will you manage the process of reverting to normal business operation when the disaster is over? That's what we mean by failback.

The good news is that this step is often less time critical and there isn't the same sense of urgency or stress as when disaster strikes.

However, it needs to be planned and the associated costs need to be assessed and accounted for.

9. Test it

Some failure procedures can be easily tested; others are far more difficult and could require serious business investment.

However, a Disaster Recovery test should be regularly planned to iron out any issues.

Some businesses plan DR tests monthly, whilst others may only do this annually. There is no right or wrong timescale on a test frequency, but choosing not to test a DR plan could prove costly during a true invocation if errors in the plan or unforeseen issues crop up that have not been properly mitigated.

It's worth pointing out that the cost of setting up a DR plan doesn't need to be excessive.

In essence, it's just the same as any other risk management process and it's something that all business should invest some time in – other than the few who don't rely on IT systems in any way.

The DR plan may be as simple as creating a documented process to invoke manual systems during the outage.

For others it may involve far more complexity such as invoking workplace recovery facilities and hosted replicated servers.

Regardless of the cost and complexity of actually implementing the DR, it's clear that not planning could have serious implications should the unthinkable happen.

And significantly increase the financial impact on the business.

The good news is that many elements of the DR plan are already in place as the result of standard IT practices and processes.

Malware is an almost ever-present threat but, without dwelling too long on IT security, a UTM, or unified threat management, solution is certainly recommended to mitigate the risk.

And it's essential to ensure that operating systems are fully patched and up-dated – something that's comes as standard with TSG SystemCare.

As businesses move towards Cloud-based services, there's an element of built-in continuity as data and systems are accessible from anywhere and on a range of devices. The priority for any business adopting Cloud-based services – other than ensuring the integrity and reliability of their cloud provider – is connectivity, so that should be a key focus for the DR plan.

Most businesses already have some backup routines in place, although there are some worrying statistics about the number who don't, and for those who have selected automated online backup solutions there's typically a relatively straightforward option to turn on DR services.

Finally, we're working with a number of customers who are moving to replication and high availability solutions, all made far easier with virtualisation, huge amounts of bandwidth and other advances in technology.

One particular customer, operating a call centre environment with a high volume of calls and all driven by information that is captured and retrieved on a live system, has reduced their RPO and RTO to zero.

Through detailed analysis, they consider even the briefest 'outage' to be a disaster and the cost implications are viewed as both unthinkable and unacceptable.

This real-time solution is built around a full hardware replication across two sites, complete redundancy and two 1 gigabyte leased line links, sourced through two separate suppliers. Crucially, the leased lines are entirely separate from the business's other connectivity requirements in order that traffic levels never compromise the integrity of these dedicated connections.

It's certainly an extreme example and most businesses won't need to invest to the same level but the principles are certainly worth applying.