

FOUR BUSINESS CONTINUITY PLANNING ESSENTIALS

Think big picture to craft an
effective business continuity plan



Introduction

Data is essential for all types of organisations today, so ensuring access to mission critical applications and data following a disaster is critical. However, business continuity and disaster preparedness are about so much more than that. In other words, you might have important apps up and running somewhere, but that doesn't matter if your office is underwater and your employees are at home without power. You need to consider the business as a whole in order to satisfy your customers needs following a disaster event.

The first step for many businesses is to conduct a business impact analysis (BIA). Detailed instruction on conducting a BIA is outside the scope of this ebook, but the point is to:

- Identify potential events that could negatively impact normal business operations,
- Calculate the likelihood that each event may occur, and
- Quantify the impact that the event could have on your business.

For example, if your data centre is near a body of water on low lying land, a flood is a possible event; its likelihood is high (during winter); and your business could be negatively impacted in a big way if downtime is significant. There are a wide variety of threats to any business ranging from natural disasters to security breaches to random accidents—a leaky pipe can have the same impact as a flood if it's directly above a critical server.

Once you have that stuff sorted, you can move on to crafting specific plans for risk mitigation, disaster response and continuity of operations. In this ebook, you will learn four distinct but interconnected business continuity planning essentials.

Crafting an employee safety and communication plan that works is absolutely essential.

1. Ensure employee well-being

Communication during and following an emergency presents a variety of challenges. So, crafting an employee safety and communication plan that works is absolutely essential. The specifics will vary widely from company to company, but your emergency safety and communication plan must address the following:

- How the company will ensure employees are safe during a disaster event; and
- How it will communicate essential information to employees following the event.

The first part will depend heavily on the nature and location of your business. Safety planning for a large manufacturing facility will obviously be very different than for a small estate agents' office, for example. Because of this, it's very difficult to provide specific best practices for this part of your BC/DR plan. However, the key is to match your safety plan to the specific needs of your organisation.

For the second part, you will need to first gather a variety of information and make sure that it is well documented, easily accessible and stored in a number of secure locations. This should include up-to-date employee contact information (email, mobile and home phone numbers, emergency contact information, etc.). It should also include a methodology for contacting employees.

Effective communication

Obviously, email is the easiest way to reach a large group of employees, but if your company's email server is down, you are out of luck. Some businesses employ redundant Exchange servers or cloud-based services to ensure email access. Of course, if you are without Internet access entirely, you'll need an alternative.

A call tree, sometimes referred to a phone tree, call list, phone chain or text chain, is another popular method for distributing important information to employees during and following an event. Here's how it works. A predetermined employee initiates the call chain with a call to the next person on the chain. That employee contacts the next person on the list and the chain continues until everyone on the call tree has been reached. Other companies may automate emergency calls with purpose-built communications software/services.



Regardless of the methods you use to distribute information to your employees, your emergency communications plan should provide enough detail that it can be carried out if the plan's creator is not available following the event (e.g. due to injury or impassable roads). Your plan should also be flexible enough to accommodate for a variety of potential emergency situations. The response to a fire in your facility during working hours will be very different from communications following the widespread distribution of a defective product, for example. Emergency communications should be brief and as accurate as possible. Depending on the structure of your organisation, you may want to keep managers updated, allowing them to pass on information to direct reports on a "need-to-know" basis. Again, the specifics of your business will dictate the correct approach.

Finally, it is essential to test and update the communications plan periodically. Testing will identify gaps in the plan such as out-of-date employee lists or contact information.

2. Keep customers in the loop

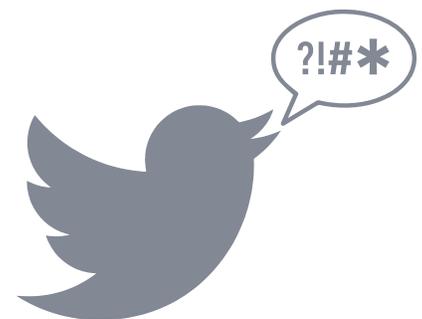
Managing customer relationships is obviously critical to the ongoing success of your business. As such, it is important to craft a plan for distributing information to your customers during and following a disaster event. The scope of your customer communications plan will vary widely depending on the nature of your business.

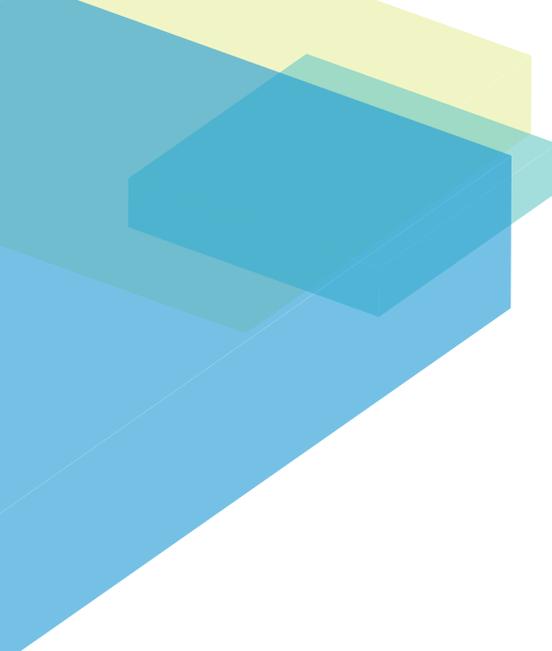
Obviously, not every glitch in operations will merit reaching out to your customers. However, if an event occurs that is likely to impact them, it is essential to communicate the details of the issue and explain the steps you are taking to mitigate it. This might mean direct communication to your customers, but it could also mean messaging via traditional and social media. Failure to do so can have a negative impact on the reputation of your organisation.

Take the way Toyota responded to reports of self-accelerating vehicles back in 2009-2010 as an example. Instead of acknowledging the issue and assuring customers that the company was investigating the problem, the company opted to cite user error in a classic example of blaming the victim. The problem was eventually pinned on floor mats, gas pedal design and faulty electronics; and although Toyota spent billions to replace accelerator components, their initial response created distrust among customers.

You will also need to handle a wide array of incoming communications following a disruption. Depending on the nature of your business this could mean: support requests, high volumes of email and phone traffic, social

Managing customer relationships is obviously critical to the ongoing success of your business.





Your organisation's ability to respond to customer needs following an event will have a direct impact on reputation.

media activity from frustrated customers, media interest—the list goes on and on. Your organisation's ability to respond to customer needs following an event will have a direct impact on reputation.

Protect your rep

So, how do you keep your good reputation intact? It comes down to careful preparation. First, you must be prepared from a personnel standpoint. Carefully planning communications with customers is essential. You will need to be able to respond quickly and clearly explain the steps you are taking to resolve issues.

All customer-facing staffers should be briefed and ready to deliver a clear and consistent message. You may want to consider using script templates, which can be adapted to address various events. Pre-scripted messages can be developed, approved by management and quickly distributed to customers following a disruption.

You also need to ensure access to communication infrastructure (phone, email, Internet access). This might mean redundant phone lines/services, hosted PBX systems, cloud-based email or redundant Exchange servers, etc. Larger businesses may need to invest in a secondary contact centre to manage inbound and outbound communications. There are a number of vendors that offer call centre services, temporary workspaces and even mobile data centres.

Testing or rehearsing all or parts of your customer communications plan should be considered essential as well. Testing is the best way to identify and resolve customer support weaknesses and communication infrastructure issues.

3. Enable IT uptime

To understand the IT piece of disaster recovery and business continuity today, it helps to look at the not-so-distant past. It really wasn't very long ago that backup meant daily incremental and weekly full backups to tape or a dedicated disk backup target. Duplicate tape copies were created and shipped offsite for disaster recovery—typically to a secondary site maintained by the business or to a tape vaulting facility (e.g. Iron Mountain). Many businesses continue to use this model today, and depending on your recovery needs it may be perfectly adequate.

However, disaster recovery from offsite tape can be painfully slow. First, you need to retrieve the tapes from an offsite location. Once they are back on premises, you must ingest data to your backup server. At that point, you can restore data and applications to your primary servers. This, of course, means considerable downtime.

**Recovery-in-place
dramatically improves
RTO because operations
can continue while primary
servers are restored**

When creating an IT disaster recovery plan, it's important to understand two concepts: recovery time objective (RTO) and recovery point objective (RPO). RTO is the amount of time that it takes to get a system restored following a failure or disaster event. So given the example above, your RTO might amount to 48 hours or more. RPO is the point in time to which data can be restored following the event. So, if you performed a backup at 6pm each night and a server failed at 5pm the following afternoon, your RPO would be 23 hours and any data created during that span would be lost. For many organisations this was unacceptable.

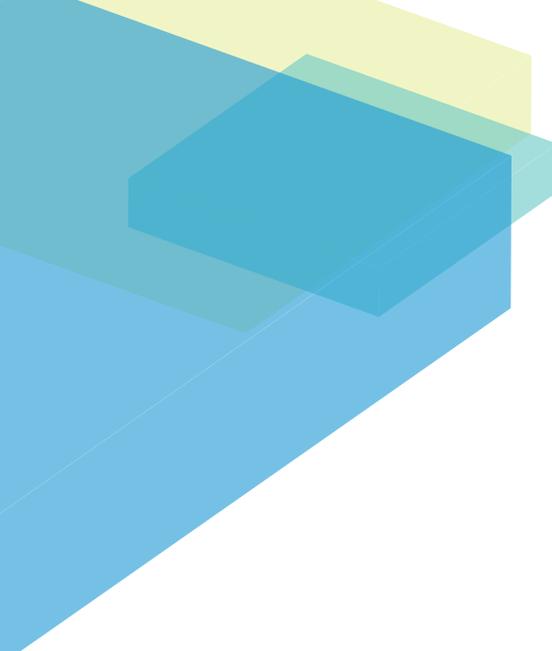
So, rather than relying on tape for disaster recovery, some organisations replicated data to a secondary site that mirrored their data centre for DR. However, this approach historically required a massive investment in hardware, because it required two sets of identical servers, storage, switches, software etc. Not to mention a secondary data centre facility. Remote replication allows users to fail over operations to a secondary site in the event of a disaster, which improves RTO, but is well out of the reach of most businesses financially.

Recovery-in-place and DRaaS

Advances in virtual server backup and cloud computing changed all of that. Today, users can run applications from image-based backups of virtual machines. This capability is commonly referred to as “recovery-in-place” or “instant recovery.” Recovery-in-place dramatically improves RTO because operations can continue while primary servers are being restored. RPO is reduced as well—snapshot-based, incremental backups at 15 minute intervals are a common practice. Virtual machine images can also be replicated to an alternate site or cloud for disaster recovery.



There are a number of ways to implement this type of system. Many backup software products today have the ability perform these tasks. If your current backup software supports it, you can set it up yourself. If you are relying on an older backup software product or you are starting from scratch, you might opt to outsource these tasks. In this model, an appliance is typically placed on premises for local backup and recovery and data is replicated to the cloud for disaster recovery. Recovery-in-place technology allows you to run applications from the onsite appliance or from the cloud following an outage or disaster. This is commonly referred to as “cloud disaster recovery” or “disaster recovery as a service” (DRaaS).



DRaaS offers the failover capabilities of traditional remote replication at a much lower price point.

DRaaS offers the failover capabilities of traditional remote replication at a much lower price point. Users typically pay a monthly subscription fee based on the amount of data they are storing in the cloud. Some services charge additional fees for the processing power necessary to run applications in the cloud during disaster recovery. Compare that with the facilities, staff and technology expenses associated with setting up a secondary data centre and the value of recovery-in-place and DRaaS is apparent.

Testing IT disaster recovery plans is essential. Historically, this was a difficult and potentially risky process. Today's technologies and services have greatly eased the testing process. Because of the ease in which virtual servers can be created, users can set up DR test environments without the risk of harming production systems. Some DRaaS providers will even perform DR testing for their clients.

4. Keep business moving

As noted above, many organisations today have limited tolerance for application downtime. If your employees or customers do not have access to essential applications and data, there will be a direct impact on productivity and revenue. While this sounds obvious, many organisations do not consider the actual costs of downtime for a business. To better understand the cost of downtime, consider the following example using Datto's [RTO calculator](#).

Let's say your business has 100 employees and on a typical day average hourly revenue is £1,500. In order to perform daily tasks, staff need access to email, a large database and a variety of file-based data. Let's say the sum of this data amounts to 2 TB and you perform an on-premises incremental backup at 6pm each day which is also copied to a cloud backup service.

Given these parameters, a full restore from a local backup would take 8 and a half hours and downtime would cost your organisation £34,000 in lost revenue. When you look at restoring 2 TB from a cloud backup following a disaster, the picture gets considerably more bleak. To restore that same 2 TB over the Internet from a cloud service it would take 6 days, 9 hours and 42 minutes and the cost to your business in lost revenue would be £614,800. Obviously, these numbers will vary widely from business to business, but this example clearly illustrates the importance of being able to continue operations while primary servers and storage are being restored.

It is critical to evaluate the facility or facilities in which your business operates.

Continuity of operations

Application downtime is, of course, just one factor that can impact your bottom line. Again, there are a broad spectrum of possible considerations depending on the size and type of your organisation. However, there are a variety of examples that apply to many businesses.

Insurance—Insurance is an important factor in your recovery effort. For example, let's say your business has numerous warehouses full of goods awaiting distribution at any given time. The cost to replace goods in the event of a fire or flood could be massive and severely impact your ability to continue operations. So, it is obviously essential to select the proper insurance coverage for your business' specific needs. Beyond that, it is also critical to document all insurance information including plan numbers/login information, the process for filing claims, etc.

Training—Every business will need to identify employees critical to the recovery process. This might mean executives, department managers and IT staff. Whatever the structure of your business, you will need to define business continuity roles and responsibilities. It is also important to cross-train staffers on essential tasks, in case a critical employee is unavailable following the event.

Facilities—It is critical to evaluate the facility or facilities in which your business operates. Considerations might include but are not limited to:

- Appropriate fire suppression systems
- Generators capable of powering essential equipment
- Uninterruptible power supply systems for critical servers
- Surge protection systems
- Alarm/intercom systems to alert employees of emergencies

Dependencies—It is important to consider dependencies within and especially outside of your organisation. Let's say you are in the business of manufacturing medical devices. You might source parts from a variety of vendors—possibly worldwide. Let's say one such vendor suffers a flood or fire and production comes to a halt. This could limit access to the raw materials you need, directly impacting your ability to continue operations. Your business continuity plan should offer solutions to mitigate these issues—for example, identifying multiple suppliers or stockpiling large numbers of essential parts.

Conclusion

Disaster recovery and business continuity planning should be considered a critical aspect of running a business. However, many organisations disregard it completely. Others have some kind of plan in place, but fail to grasp how time consuming the recovery process can be and the the associated cost of downtime. The good news is that today's data protection technologies and services have greatly improved the IT piece of the business continuity puzzle. There are a wide array of options in the market today at different price points, which enables you to select a product or service tailored to your specific business needs.

As you may have noticed, testing your plans has come up throughout this ebook. The importance of testing business continuity/disaster recovery plans can not be understated. Testing is the only way to reveal gaps in your plans and address them proactively—not while you are frantically trying to pull the pieces back together after heavy rains deposited a foot of water in your reception.