wanstor

Disaster Recovery Planning: The Essentials

A guide for IT Professionals

Contents

Introduction

For many business or not-for-profit organisations, an IT 'disaster' can often refer to something impacting their data from outside, such as storms, an earthquake, acts of vandalism or sabotage.

While major events like these should lead IT to reflect on its preparations, disaster recovery assessment must not be limited to the consequences of a flood, a fire or similar natural disaster.

Lower profile yet still important events - from software bugs to hardware failures - can prove just as (if not more) catastrophic as fire, power outages or flooding, and must be given due consideration as well.

At Wanstor we believe an IT disaster is not limited to what makes the news or captures the leadership team's attention, but rather anything that makes the ordinary conduct of business either difficult or impossible.

If an event - at any scale - can interrupt IT operations, it poses a threat that must not be ignored. Whatever is at stake, be it the loss of revenue, reputation or customers, represents a potential disaster which the IT function must either be prepared to avoid or from which it must be prepared to recover.

In this white paper, Wanstor experts offer a business perspective on what can mistakenly be considered a technology issue.

The most crucial considerations should be determined by business needs over IT requirements. In fact, the most important disaster recovery decisions are not around technology, but around business demands that drive technology choices.

Technologies for data recovery and application availability have evolved significantly over the past five years, but the underlying business reasoning, core to any effective disaster recovery plan, remains consistent and should include considerations for assessing business exposure to disaster, a review of options for cost-effective preparation and recovery, setting expectations for performance that direct technology decisions, and testing disaster recovery plans for vulnerabilities.

The first step towards disaster recovery planning is awareness: understanding what a disruption would mean to your business, and what you can do to prevent or mitigate disastrous consequences.

Assess your business needs

If you were to ask the IT team to assess your vulnerabilities, the chances are you would receive a reasonably accurate account of your data and applications with pages of documentation around redundant drives, backups, and remote data centres.

This type of report is likely to expose the technological consequences of a disaster. It would however not reveal the business consequences of lost hours, lost data and lost applications, leading to lost revenue, profits, customer confidence, or worse.

To expose these consequences, IT directors should conduct a business impact analysis, or BIA. A BIA calculates the monetary loss of a single event on the business they are employed by - such as a fire, hardware failure, flooding or software bug.

This takes into account the extent of any damage - for example how much data has been lost or how many interactions or integrations have been interrupted - and the duration of these disruptions, as well as how long it takes to restore said data, applications or operations.

Using reporting of this kind, IT directors will be able to calculate a number that represents potential loss – the quantifiable sum of everything that may be at risk in the event of sudden disaster.

> **From a business perspective, an IT disaster isn't just what makes the news or captures attention at management level, but rather anything that makes day-to-day business either difficult or impossible**

This report may then be used to inform the wider organisation of these consequences, and the potential impact of disaster in their respective business areas.

They can also advise on precautionary measures that need to be taken in order to prevent IT disasters from leaving business critical assets, data and processes at risk.

Are you missing silent disasters?

In practice it is easier and usually quicker to recognise an immediate disaster that would result in losses at scale.

This is compared to the recognition of a small or silent disaster that may take weeks or even months to come to the fore with damage already been done across the IT estate.

While most IT professionals intuitively understand the consequences of a loss at scale, most will fail to recognise the extent of a silent IT disaster unfolding on their watch.

According to IT complexity expert and ObjectWatch founder Roger Sessions, organisations across the world lose £6.2 trillion from IT failures every year. Although Sessions' numbers have been challenged by other economists, these calculations remain sobering.

The most notable aspect of Session's math is the overwhelming majority of annual losses not being caused by the low probability or high-consequence catastrophes that capture attention.

They are predominantly caused by high-probability and low-consequence events that occur frequently, such as software bugs, hardware failures and security breaches.

Even worse is that as applications become more complex, involving larger swathes of code, data nodes and systems networks, the exposure to smaller events becomes more frequent and their impact more costly.

If your only assessment of loss is a business impact analysis, IT directors may be missing the real cost of IT disasters and failing to plan for recovery.

> While it is still important to conduct a business impact analysis, IT decision makers must not allow results to blind them to the consequences of multiple high-probability events that, year after year, will impose losses on the organisation.

These professionals need to recognise that potential loss from a catastrophic event must be complemented with a deep understanding of expected loss - a more realistic figure that factors in two critical elements.

Probability: Expected loss includes calculations for probability, the likelihood of a loss event that potential loss neglects. When probability is accounted for, the significance of multiple, small events becomes visible, allowing IT directors to direct budgets and resources to the events that really matter: the ones that regularly impact the bottom-line.

Current investments: Quite often, the business impact analysis fails to account for current investments in recovery - such as backups and automated failovers - that would temper overall losses. Although the process may seem intimidating, there's a better outcome that can be achieved with this calculation as IT directors can make a more sophisticated loss assessment.

No matter how much the IT team has invested in preventative or corrective action, potential loss - because it measures the overall value of business at risk - never actually decreases.

When an IT director shifts focus to expected losses, which include accommodation for probability and corrective action, they can actually see a reduction in loss exposure - and truly measure the value of disaster recovery investments.

Going beyond business impact analysis

Action steps

1. Include probability into risk calculations to arrive at realistic expected loss figures.

2. Shift focus so that high-probability and low-consequence events figure as or more prominently in disaster recovery planning than low-probability and high-consequence catastrophes.

3. Identify and protect hidden dependencies (such as supplier networks, access to physical buildings or even availability of personnel during a disaster) that must be taken into account to recover critical data and applications.

4. Establish priorities. Not all data and applications are equal; the bulk of disaster recovery planning should be directed toward the top 20% of expected losses.

Review Your Options

Not all applications and data are equal, in view of their business impact. Some require greater investment in disaster recovery while for others, lower standards for recovery may be tolerated due to the fact that were these to be affected, the impact would be less should disaster strike.

Technology choices should mirror business objectives, with priorities established in the previous section dictating the level of investment made in disaster recovery.

Two key decisions: Recovery Point and Recovery Time Objective

At the heart of any disaster recovery plan are two critical decisions that reflect an IT team and business tolerance for loss:

+ The recovery point objective (RPO) that determines the moment in time, before the disrupting incident, that is restored to. The closer the RPO to the incident, the lower the data loss.

+ The recovery time objective (RTO) that establishes the amount of time it takes to restore operations. The lower the RTO, the less time it takes to recover.

Match Service Level Agreements to Priority Tiers

RPO/RTO of Seconds to Minutes: This category includes data and applications which are that important - measures of public safety (health, military, police) or financial impact (banking, insurance, trading) - that they demand a zero RPO and a zero or near zero RTO.

Meeting your obligations will require investments in automated solutions that can respond instantly to disaster.

RPO/RTO of Minutes to Hours: Here, data and applications are important but not mission-critical. IT teams should think Enterprise Resource Planning, Customer Relationship Management and email for example.

Automation still plays a role, but you can accept some minor data loss from your RPO, and can endure a few hours of recovery time delay.

RPO/RTO of Hours to Days: Consider this the place for less critical but *nice to have* applications, such as the intranet or human resource functions.

Time is not of the essence, and much of the disaster recovery can be managed through inexpensive manual efforts.

Action steps:

1. Assess, rank and categorise your data and applications by their business or safety significance.
2. Assign different RPO and RTO performance requirements to different tiers.
3. Budget unequally and always anticipate a higher spend on most critical tiers.

By definition, the IT team has technology expertise, but as the business decision-maker, the IT Director must set the objectives the technology must achieve. Highly ranked amongst these is simplicity, when disaster strikes, recovery must be simple and easy if it is to be rapid and effective.

Fast, easy recovery requires consideration of the following:

Automation: In the event of a disruption or emergency, the IT director will not have time to assemble teams, coordinate meetings and distribute responsibilities.

To meet previously determined RPOs and RTOs, IT directors need events-driven application management, an automated process that eliminates or minimizes manual intervention.

Comprehensive fit: Your organisational IT infrastructure wasn't built in a day, but will have taken shape over years, incorporating a mix of environments (physical or virtual), platforms and operating systems.

Regardless, disaster recovery technology must work across all components to communicate and coordinate events among disparate pieces.

Availability and reliability: Data storage resilience – the ability to recover quickly from failure – must be accompanied with data that makes sure all systems that require the recovered data can find and access it.

Simple restoration of complex applications: That one purchase on an ecommerce site or that one withdrawal from an ATM? Behind the scenes, these single activities represent a complex, multi-tiered stack of technology that often includes application code, stored data access, middleware connectivity, and other functional layers.

Effective recovery requires technology to restore each layer, in the right order, and then reintegrate their activities to recover an entire application. If current recovery technologies cannot restore the entire stack of a multi-tiered application, they cannot perform the most business-critical technology functions.

Action steps:

1. Assess your current disaster recovery components. Are they integrated and automated for rapid action, or will your recovery be delayed by the need for coordinated manual interventions?

2. Review application layers to be sure that every tier can and will be restored, in the right order, in the event of disaster.

3. Conduct an IT inventory to expose the system elements and dependencies that must be restored together to effect a rapid recovery.

Considering Recovery Capacity Objectives

Recovery objectives are the mainstay of your disaster recovery plan. But if you are obligated to fulfil service level agreements (SLAs) for customers, you should consider a third metric: the recovery capacity objective (RCO).

This is the acceptable amount of functionality you require - not only to recover, but also return to the contracted level of service you are obliged to reach. Your RCO represents performance that can vary from a compromised level of restoration to a full service return.

A dive into Data Replication

Data replication refers to the process by which data in one site is mirrored in another, typically the backup location designated for disaster recovery.

There are different types of data replication, each with its strengths and weaknesses:

Synchronous: With synchronous data replication, data tasks at the primary site are not acknowledged as completed until they have been replicated at the secondary site.

While synchronous replication closes the RPO gap, it comes with some drawbacks in system performance, and its application is limited to sites within a wide area network typically no more than 60 miles apart.

Asynchronous: Asynchronous replication accommodates mirroring data across any distance and allows the primary centre to write to disc without waiting for acknowledgment from the secondary.

Although asynchronous replication allows for greater speed and distance, it opens up a gap in the data record between the two sites, potentially compromising RPO.

Hybrid: The hybrid approach applies both methods, using synchronous replication for almost instantaneous availability in the event of localised failures, and asynchronous replication to a distant data centre to provide restoration in the event of a disaster.

The client determines the failover threshold from one site to the other, and the application of the hybrid solution requires sophisticated planning.

Automatic vs Automated: What's the difference?

Automatic requires no manual intervention whatsoever; the triggering event initiates a sequence of activities almost instantaneously.

Automated refers to processes that, once initiated by a manual action or decision, run without further need for intervention. For localised recovery, such as disc to disc or even a failover to a nearby data centre within a wide area network, automatic solutions are preferred.

But for failovers to distant data centres that might impose disruptions to data streams, automated processes give businesses the power to make informed choices.

Test Your Plan

At this point the IT director has assessed their needs, established priorities, matched service levels to those priorities, and set expectations for recovery solutions.

Once the appropriate technology has been identified and purchased, you may expect that IT would then be prepared for disaster recovery, but this is not yet the case.

The final step, one which IT need to take in ensuring their plan truly meets business needs, is one they will need to repeat time and again - testing.

You do not test to prove that your plan works; you test to expose vulnerabilities and make the unknown known, *before* disaster strikes.
If as an IT function you have never failed a disaster recovery test, it can be said that you do not have a comprehensive disaster recovery plan in place.

By actively searching for and finding holes in the IT disaster recovery plan, IT leads can make informed business decisions.

These decisions may include such that if the probability of a particular failure is low or the consequences of said failure are minor, it may be decided that additional protection is not worth the added expense.

If, however, vulnerabilities are deemed to be highly probable, or could have significant consequences - or both - IT professionals then know precisely where to direct disaster recovery investments.

Use the following checklist to make sure you have determined the recovery needs and technology objectives your IT team must execute effectively.

+ Have you linked IT functions to business consequences and assigned a monetary value to their significance?
+ Does your definition of disaster include the high-probability and low-consequence events that cause most catastrophic business disruptions?
+ Can you calculate, not just potential loss, but expected losses?
+ Do your calculations reflect both current mitigations and event probabilities?
+ Have you used your expected loss figures to focus your disaster recovery priorities?
+ Do your RPO and RTO service levels reflect your priorities?

- Have you created a hierarchy of tiers that allow you to make recovery investments matched to the business significance of your applications and data?
- In addition to your RPO and RTO, have you set a recovery capacity objective (RCO) that acknowledges gradations in recovery status?
- Do you have the appropriate data replication model for your recovery needs?
- Are your recovery solutions automated to facilitate rapid, coordinated recovery in the event of disaster?
- Can your current recovery solution embrace your entire technology environment, regardless of platform, operating system, and other variables?
- Will your stored data be ubiquitous upon restoration, available to every application and system that needs it?
- Can your recovery solutions restore every layer in your complex, multi-tiered applications, automatically and in the correct order?
- Do you regularly test your disaster recovery plan, not to prove efficacy, but to expose vulnerabilities?

If you cannot answer yes to each of these questions, you have areas in your disaster recovery plan that may require closer attention.

Wanstor Managed Backup & Disaster Recovery Services

Wanstor's Disaster Recovery Service provides businesses with regular replication of critical applications, infrastructure, data and systems for rapid recovery after an IT outage.

Disaster recovery is a critical IT feature that every business, large or small, should employ. Without it, thousands if not millions of pounds could be lost or an entire business reputation wiped out if critical IT systems are not backed up or cannot be recovered quickly.

From hardware failures to large-scale natural disasters, IT teams must be prepared for when a disaster happens and have the ability to get the IT operation up and running again in the shortest time possible with minimal inconvenience to customers and / or business users.

The key to a high-performing IT disaster recovery plan is having the right mix of solutions to achieve your businesses need for speedy recovery and maximum value.

Wanstor has architected a suite of Disaster Recovery-as-a-Service solutions to help businesses achieve their goals around IT service availability and data protection.

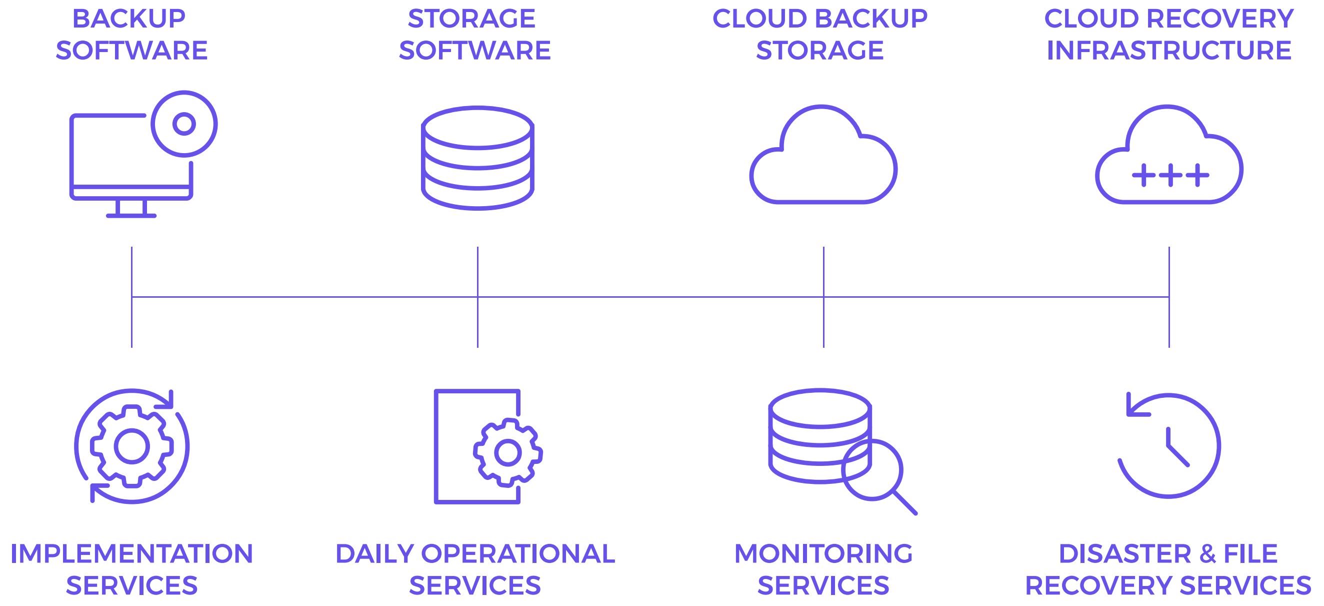

Figure 3: Wanstor Backup and Recovery Services

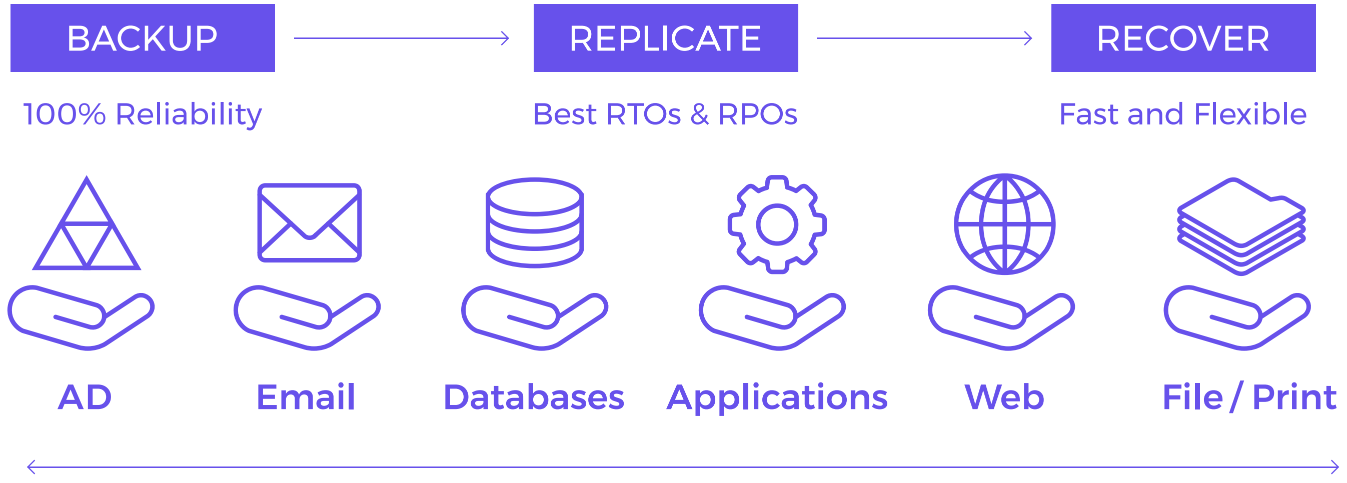

Figure 4: Wanstor's Backup and Replication Services

Wanstor Backup & Disaster Recovery Services

Business continuity solutions help maintain employee productivity and support your business's ability to generate revenue.

When businesses experience downtime for any reason, they cannot function as usual. This in turn means lost revenue.

A comprehensive backup and disaster recovery solution ensures that your organisation is operational as soon as possible after a disaster.

Backup and disaster recovery solutions help preserve a company's reputation with customers and partners. Downtime leads to more than only financial drain on the business - the long term reputational costs of may prove disastrous.

Business continuity solutions can also help to prevent companies losing business to competition. The more competitive pressure your business is under, the more downtime may lead customers to defect to competitors.

A sound backup and DR solution can help reduce customer churn by ensuring systems are always on and your clients are not inconvenienced in any way.

Backup and disaster recovery solutions can also help to ensure compliance with industry regulations. A backup and disaster recovery plan ensures businesses do not have to worry about compliance violations and legal issues related to data loss and downtime.

Disaster Recovery services that Wanstor offers include:

+ Assessing DR and BCP requirements alongside your IT that takes into account budgetary decisions and business impact
+ Architecting solutions to support RPO and RTO objectives
+ Replication of data and systems with backup available in Wanstor's own private cloud
+ Implementation of failover and replication at network, application and storage layers
+ Creation, testing and maintenance of DR plans
+ The ability to consult on, deploy and manage Disaster Recovery Service solutions

Why Wanstor?

A Managed IT Service Provider in London, Wanstor brings industry leading expertise and capability.

Along with meeting criteria in each area of consideration when selecting a managed IT service provider, Wanstor has experience in helping customers manage diverse and complex IT environments.

We have the technology and business knowledge to help you understand and identify your IT and business requirements at every stage in the IT Managed Service journey.

We specialise in each type of delivery model, including backup and disaster recovery services, cloud, traditional IT and strategic outsourcing, helping you to realise an integrated multi-sourcing IT strategy.

We offer flexibility to select the degree of support you want for each layer of infrastructure – from basic monitoring and management, to designing and delivering a roadmap to digital transformation.

For more information on how Wanstor can help your business and IT department achieve its objectives, request a call back at the link below.

Request a call back

wanstor