

A photograph of a man and a woman in an office environment. The man, wearing a mustard yellow shirt, is standing and gesturing with his hand while looking at a computer monitor. The woman, wearing a teal shirt and white earbuds, is sitting at a desk and smiling. The background is a blurred office space with other monitors and people.

wanstor

Disaster Recovery: A Guide for IT Professionals

White Paper

Contents

- + Introduction
- + Writing the Disaster Recovery plan: What you need to know
- + Developing your Disaster Recovery strategy
- + Make Disaster Recovery strategies a part of the overall plan
- + Developing the Disaster Recovery plan
- + Incident response
- + Other activities
- + Virtual Disaster Recovery
- + Risk reduction
- + Why server virtualisation for Disaster Recovery?
- + Cloud-based Disaster Recovery
- + Defining the Disaster Recovery blueprint with cloud computing
- + Cloud recovery options
- + Disaster Recovery best practice tips
- + Final thoughts

Introduction

Disaster Recovery and Business Continuity as an IT topic have moved up the leaderboard of concerns for IT directors over the past few years.

At Wanstor, we see many organisations discussing Disaster Recovery plans. But how feasible are these plans, should the need for deployment ever arise?

Quite often our IT experts find that Disaster Recovery plans stand up under high level scrutiny, but when interrogated closely, are found lacking further information or data pertaining to those who may access further information.

With great strides made in Disaster Recovery over the last five years, there is no excuse for businesses not to have a basic Disaster Recovery plan in place should disaster strike. Even those new to IT management can learn from the finance industry around Disaster Recovery, with this particular market sector having undergone radical transformation to offer secure, *'always on'* systems.

Regulations such as Basel II, MiFID and those of the UK Financial Services Authority have helped dictate the standard of Disaster Recovery expected, and the minimum distance of secondary sites.

The basics of DR stipulated under this legislation provide most businesses, no matter what industry sector, the tenements of a Disaster Recovery plan. In summary, businesses without these plans leave themselves exposed to reputational damage and restrictions on trade.

It has been proven over time that those experiencing disaster without a recovery strategy in place often fail to recover; customers lose confidence in the brand, and loyalty fades.

While there are compelling push factors when considering marketplace and customer confidence, there are attractive pull factors making Disaster Recovery simpler, more cost-effective and easy to manage. Server virtualisation, cloud computing and the wealth of access to data centre co-location providers mean that developing and implementing a Disaster Recovery plan is easier now than at any time in the past.

In this white paper, Wanstor's Disaster Recovery experts outline what you, the IT Director, need to know about developing and implementing a Disaster Recovery plan.

Writing the Disaster Recovery plan: What you need to know

Developing a detailed Disaster Recovery plan is the overarching aim of any IT DR planning project.

It is through these plans that IT Directors will develop the detailed steps required in recovering IT systems to a state where they can support the business after a disaster.

Before any plan takes shape, the IT director must perform a risk assessment (RA) and / or a business impact analysis (BIA) so as to identify those IT services that support the businesses critical operational activities.

They will then need to establish recovery time objectives (RTOs) and recovery point objectives (RPOs). Once this work is complete, IT Directors will be ready to move forward with developing Disaster Recovery strategies, followed by the actual plan.

In this section of the document we will explain how to develop Disaster Recovery strategies, as well as how to write a Disaster Recovery plan.



Developing your Disaster Recovery strategy

In terms of Disaster Recovery strategies, ISO / IEC 27031 is the global standard for IT Disaster Recovery.

It states that ‘Strategies should define the approaches to implement the required resilience so that the principles of incident prevention, detection, response, recovery and restoration are put in place’.

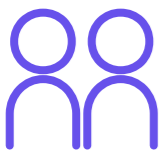
Strategies define what you plan to do when responding to an incident, while plans describe how you will do it.

Once you have identified your critical systems, RTOs and RPOs, we suggest that you create a table (Figure 1 opposite), helping with formulation of the Disaster Recovery strategies you will need to protect each department and their relevant IT systems.

As part of a Disaster Recovery plan, IT Directors should consider not only the affect on account systems, but also issues around budgets, the position of senior management on business risk, availability of human, financial and systems resources, costs versus benefits, human and technological constraints, and regulatory obligations.

Critical System	Accounts Payable	Manufacturing	Building Security
RTO / RPO (in hours)	4/2	8/4	2/2
Threat	Server Failure	Loss of manufacturing systems	Security System destroyed
Prevention Strategy	Secure equipment room and backup server, install UPS	Set up failure alerts and conduct regular inspections, install UPS	Move to secure area, adding protective enclosures around sensor units, install UPS
Response Strategy	Switch over to backup server, validate that UPS is running	Run manufacturing on alternate system	Deploy guards at strategic points
Recovery Strategy	Fix or replace primary server, fall back to primary server	Fix primary manufacturing system, resume normal operation	Obtain and install replacement units and sensors

Let’s take a closer look at some additional factors in strategy definition.



People

Includes availability of staff and contractors, training needs of staff and contractors, duplication of critical skills to primary and backup personnel, available documentation for personnel, ensuring retention of knowledge.



Physical Facilities

Availability of alternate work areas on-site, at one alternative business site, one third-party-provided site, employees’ homes, or at a transportable work facility. Consider site security, staff access procedures, ID documentation, location of alternate sites relative to primary site.



Technology

Consider access to equipment space properly configured for IT systems – raised floors, suitable heating and ventilation, electrical power, voice and data infrastructure, distance between alternate technology area and primary site, failover and failback technologies.



Data

Timely backup of critical data to secure storage area in accordance with RTO / RPO requirements, method(s) of data storage, connectivity and bandwidth requirements so all critical data is backed up according to RTO / RPO timescales, data protection capabilities at alternate storage site, availability of technical support from qualified third-party service providers.



Suppliers

Identify process for nominating and contracting suppliers for all critical systems and sourcing of people. Key areas of alternative supplier importance include hardware (servers, racks), power (batteries, universal power supplies, power protection), networks (voice and data network services), repair and replacement of components, multiple delivery firms (FedEx, UPS).



Policies and Procedures

With Disaster Recovery policies defined, gain approval from senior management to sense check everything required under the plan. Define step-by-step procedures to initiate data backups to secure alternative locations, relocate operations to alternative spaces, recover systems and data at alternative sites, and resume operations at original site or a new location.

Make DR strategies a part of your plan

Once your Disaster Recovery strategies have been developed, time should be taken to translate these into Disaster Recovery plans.

Figure 2 opposite illustrates critical systems and associated threats, response strategy and new response action steps, as well as recovery strategy and new recovery action steps.

This approach can help IT Directors to quickly drill down and define high-level action that should be taken.

From this table you can then expand the high-level steps into more detailed procedures, as necessary; always ensure, however, that these are sequenced correctly.

This will save time should a disaster actually take place.

Critical System	Accounts Payable	Manufacturing	Building Security
Threat	Server Failure	Loss of manufacturing systems	Security System destroyed
Response Strategy	Switch over to backup server, validate that UPS is running	Run manufacturing on alternate system	Deploy guards at strategic points
Response Action Steps	Verify server is down, verify data is backed up, test backup server, switchover to alternative server	Verify system is down, verify data backed up, test alternate system, switchover thereto	Verify security system is down & data is backed up, source, brief and equip on-site guards
Recovery Strategy	Fix or replace primary server, fall back to primary server	Fix primary manufacturing system, return to normal operation	Obtain and install replacement units and sensors
Recovery Action Steps	Verify server outage cause, obtain, test & install new server, fail systems back to new server	Verify cause of outage, contact repair resource, fix & test system, fail back to repaired system	Verify cause of security system outage, replace, test and restart security system

Developing the Disaster Recovery plan

Disaster Recovery plans should provide a step-by-step process covering your response to a disruptive event.

Procedures in place should ensure an easy-to-use, repeatable process for recovering damaged IT assets and returning them to normal operation as quickly as possible.

If staff relocation to a third-party site or other alternate space is necessary, procedures must be developed for these activities at the same time.

Incident response

As well as strategies developed under a Disaster Recovery plan, IT Disaster Recovery plans should form part of an incident response process that addresses initial stages of the incident and steps to be taken.

This process can be seen as a timeline, such as a *Disaster Timeline*, in which incident response actions precede Disaster Recovery actions.

We believe the best Disaster Recovery plans should begin with key action steps and list key contacts in order to simplify authorisation and launch of the plan should disaster strike.

The following page illustrates a suggested plan structure based on Wanstor's best practice knowledge gained through working with hundreds of businesses across the UK when implementing their IT Disaster Recovery plans.

Introduction

Following initial emergency pages, DR plans should have an introduction that includes the purpose and scope of the plan. This section should specify who has approved the plan, who is authorised to activate it and a list of linkages to other relevant plans and documents.

Document History

A section on plan document dates and revisions is essential, including revision dates, details and information on those who implemented said revisions. This should be located at the front of any plan document.

Role & Obligations

The next section should define roles and responsibilities of DR recovery team members, their contact details, budget spend limits, secondary signatures, and the limits of their authority in a disaster situation.

Incident Response

The incident response process alerts IT to abnormal incidents (e.g. system alarms being triggered), allowing situation (and damage) assessment for early determination of severity, incident containment to restore control, and notification of senior management along with other key stakeholders.

Plan Activation

Based on findings from incident response activities, the next step is to determine if Disaster Recovery plans should be launched, and which ones in particular should be initiated. Once plans are initiated, incident response activities can be scaled back or terminated, depending on the incident. This section covering activation should define criteria for launching DR plan/s, what data is needed, and who makes this determination. Included within this part of the plan should be assembly areas for staff, procedures for notifying and activating DR team members and procedures for standing down if management determines this form of response is not required and incidents can be dealt with via alternative means.

Procedures

Once the plan has been launched, DR teams should take assigned materials and proceed with response and recovery activities as specified. The more detailed a DR plan is, the more likely that your affected IT assets may be recovered and returned to normal operation. Technology DR plans can be enhanced with relevant recovery information and procedures obtained from system vendors. Also check with these vendors while developing your DR plans to see what they will provide in terms of specific emergency recovery documentation.

Appendixes

Located at the end of a DR plan, appendixes may include systems inventories, application inventories, network asset inventories, contracts and service-level agreements, supplier contact data and any additional documentation that may facilitate recovery.

Other activities

Once your DR plans are complete, they are ready to be exercised. This process will determine a plan's capacity for recovery and restoration of IT assets as forecast.

Included alongside this are processes including employee awareness, personnel training and records management.

These are essential in ensuring employees are fully aware of DR plans and individual responsibilities during a disaster and that DR team members have been adequately trained in specific roles as defined within your plan.

Records Management and Change Management activities should be initiated as part of any Disaster Recovery plan

Since DR planning generates a significant amount of documentation, both records management (and change management) activities should also be initiated.

If your business already has both records and change management programmes, these should be utilised within your DR planning.

Virtual Disaster Recovery

Storage and server virtualisation make many onerous Disaster Recovery tasks easy to execute whilst reducing overall DR costs.

If your business still lacks a viable Disaster Recovery strategy, give serious consideration to virtualisation. Reasons why the adoption of server virtualisation has become so popular as part of business Disaster Recovery plans are that it can improve resource utilisation and lower IT costs through consolidation, whilst also improving system availability.

In summary, virtualisation turns physical devices into resource pools independent of the physical assets which they run on. With server virtualisation, the decoupling of operating systems, applications and data from specific physical assets helps to eliminate economic and operational issues of infrastructure silos – a key factor in having an affordable Disaster Recovery strategy.

Storage virtualisation takes those same benefits and extends them from servers to the underlying storage domain, bringing IT teams closer to the ideal of a virtualised IT infrastructure. By harnessing the power of virtualisation at both the server and storage level, IT teams become more agile in Disaster Recovery.

Risk reduction

Senior Management in many industries are becoming aware that an IT disaster could have a terminal impact on operations.

This leaves IT Directors constantly challenged in ensuring Disaster Recovery plans are in place for their organisations.

While most businesses may execute daily data protection plans, few focus their efforts on true Disaster Recovery strategy, such as event interrupting service at any primary production location. Such incidents take many shapes, including power failures, fires, floods, weather related outages, natural disasters or terrorism.

Regardless of cause, unplanned downtime in the data centre causes major headaches around IT's ability to maintain business operations. The overarching goal of any DR strategy should be to recreate all necessary systems at a second location as quickly and reliably as possible.

Unfortunately, DR strategies often comprise fragmented thinking, with staff involved in generating DR documentation finding themselves moved to other projects and programmes of work. Strategies composed under such circumstances are at risk of becoming disjointed, costly and overly complex.

This approach to Disaster Recovery planning, alongside the belief that existing backup processes are adequate, means that some businesses are not prepared when disaster strikes.

At Wanstor we know that process and backup technologies will only take a business so far when it comes to Disaster Recovery.

When relying on backup for DR, the time taken to acquire replacement hardware, reinstall operating systems and applications and recover data – even from a disk-based copy – will likely exceed a recovery time objective (RTO) of one to three hours, under the most favourable conditions.

Recovery from a mirror copy of a system is faster than recovery with traditional backup methods, but it is also more expensive and complex. Maintaining identical systems in two locations and syncing configuration settings and data copies can be a challenge.

This may prompt a business to prioritise their data, providing greater protection to specific tiers over others. The most critical type of data is Tier 1 data, representing roughly 50% of any organisation's total data.

Why server virtualisation for Disaster Recovery?

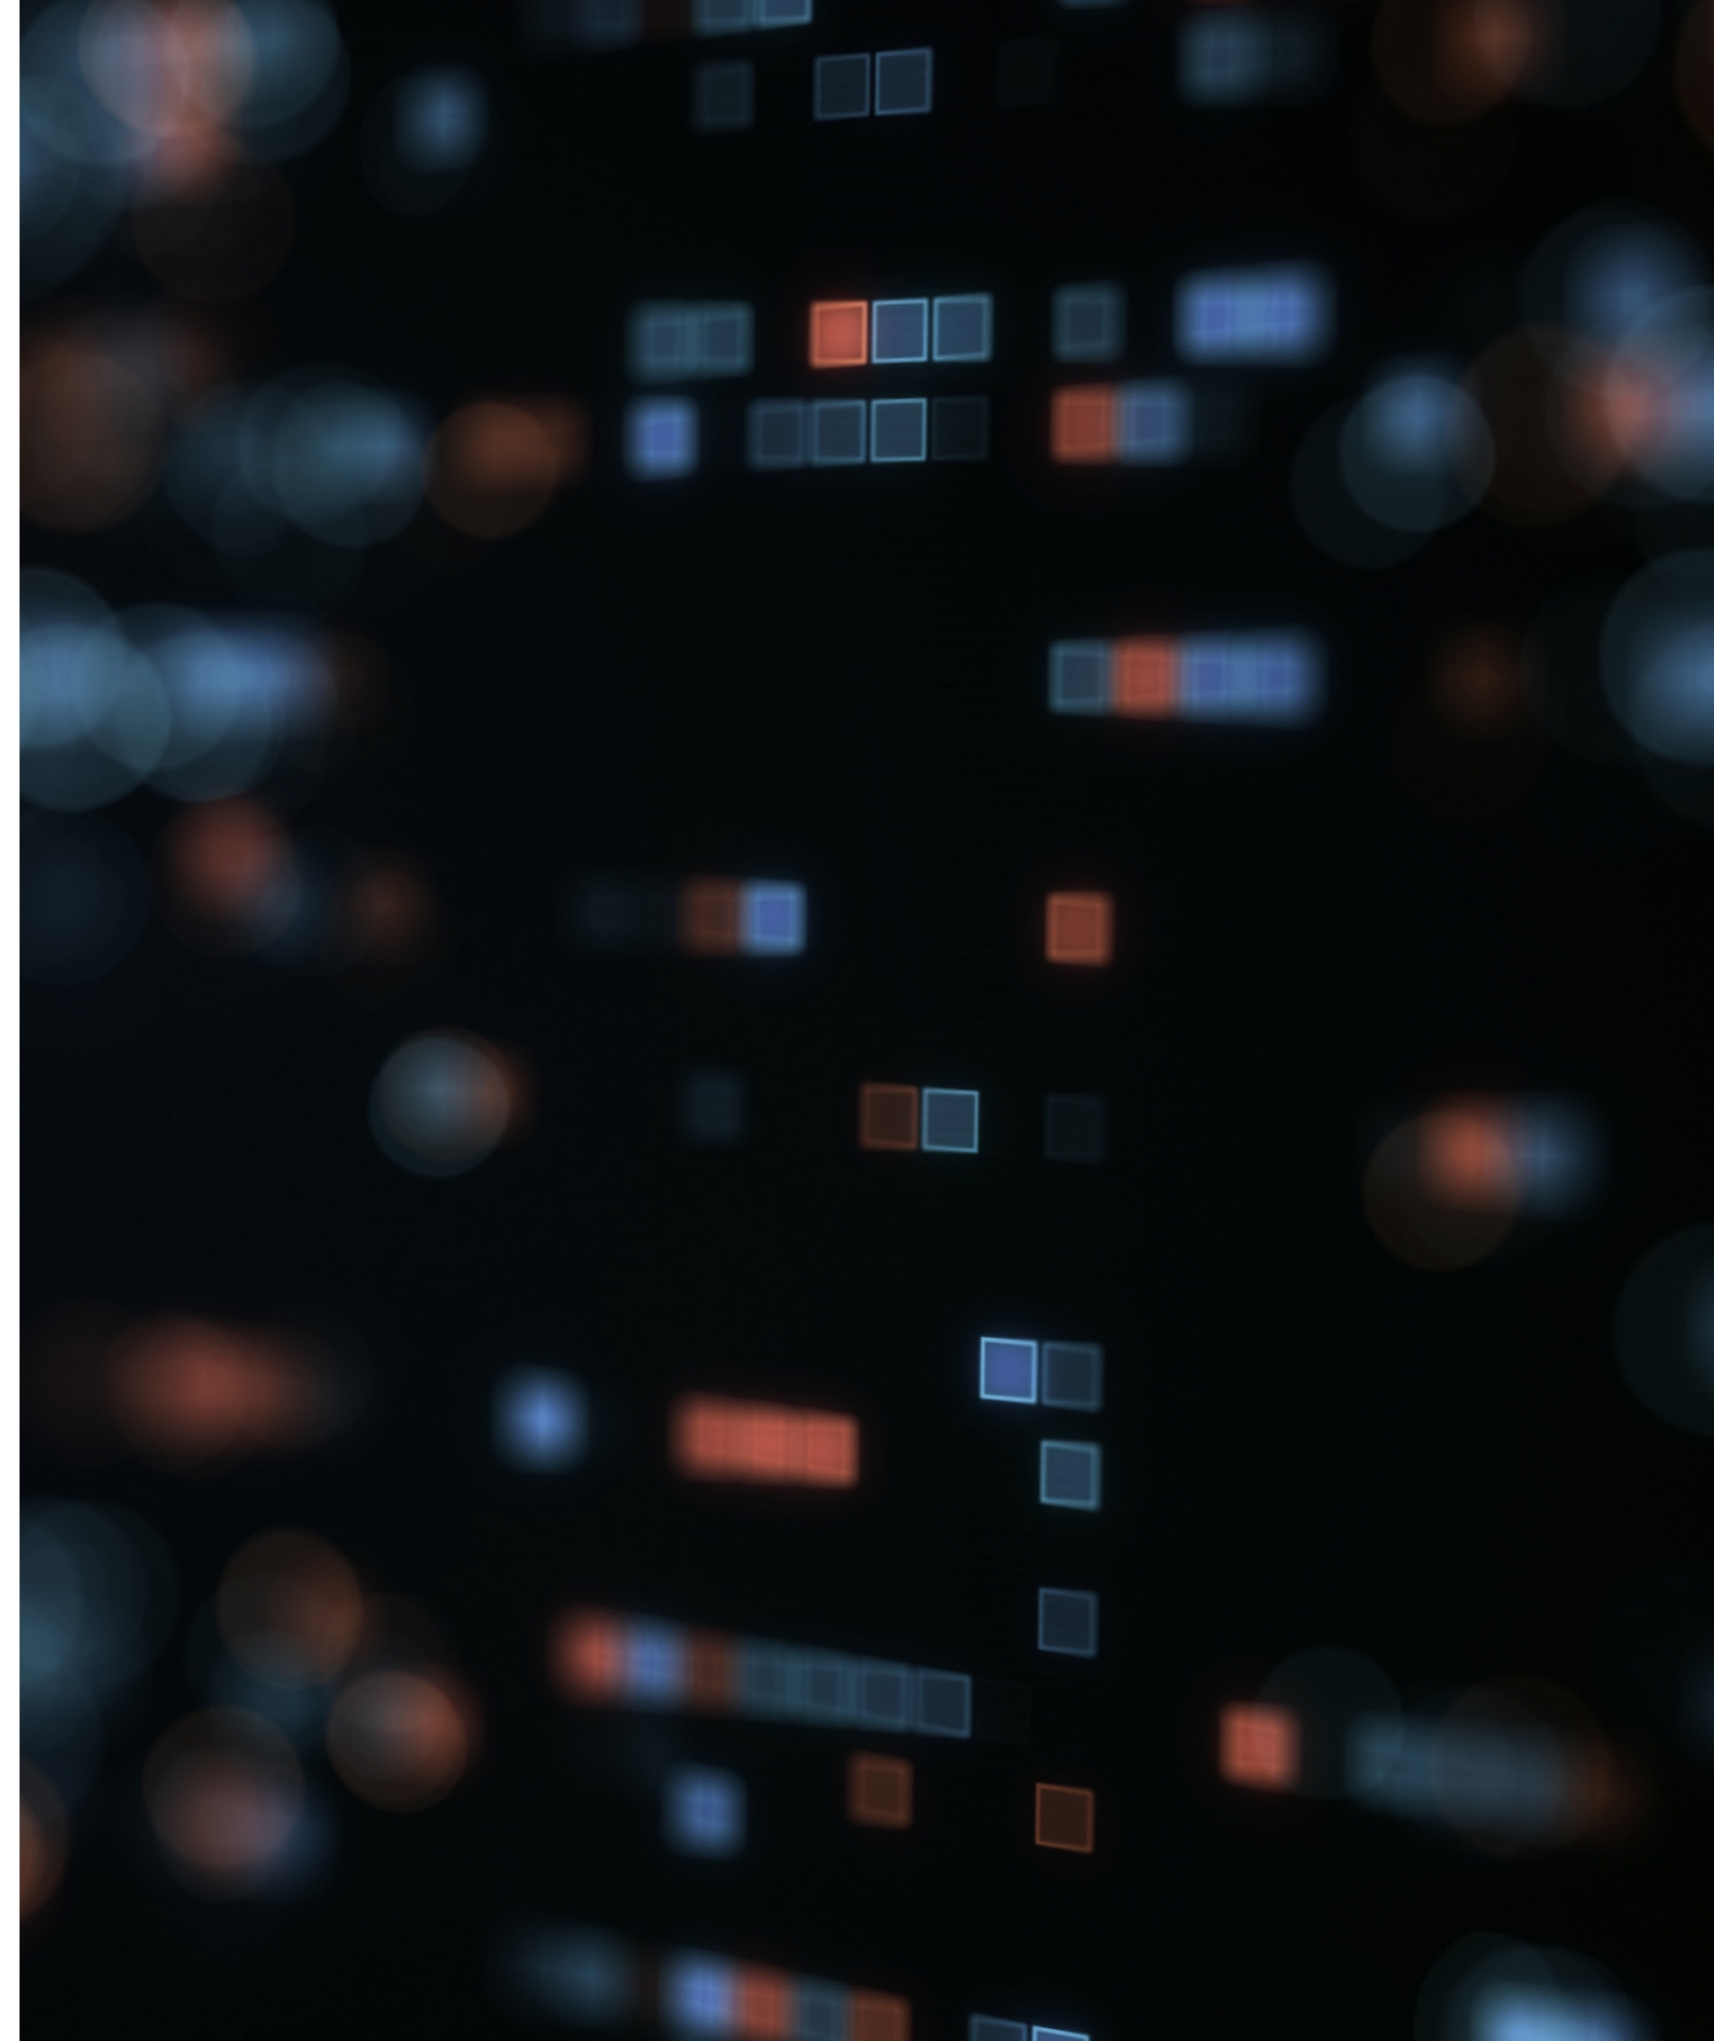
Virtualisation has become a major catalyst for change in x86 environments, as it provides new opportunities for more cost-effective DR.

When examining the reasons behind server virtualisation, most IT Directors have told us that using virtual machine replication to facilitate Disaster Recovery is a primary reason for selecting server virtualisation as part of any Disaster Recovery strategy.

This is because server virtualisation abstracts from the physical hardware layer, eliminating the need for identical hardware configurations at production and recovery data centres.

Since virtualisation is quite often a catalyst to refreshing underlying infrastructure, there may be retired hardware available.

For businesses who are not able to secure the CapEx for a DR configuration, they may as a result be able to leverage old, retired or used hardware.



By consolidating multiple applications on a single physical server at recovery data centres, physical recovery infrastructure required is reduced. This in turn minimises expensive raised floor space costs, as well as additional power and cooling requirements.

Leveraging the encapsulation and portability features of virtual servers aids in DR enablement. Encapsulating a server into a virtual machine enables mobility and allows multiple copies of virtual machines to be created and easily transferred within and between sites for business resilience and DR purposes.

This offers a dramatic improvement over backup of data to portable media (such as tape) or rotating media at a cold standby site.

Additionally, protecting virtual machine images and capturing the system state of virtual machines are relatively new concepts not previously available when considering physical hardware.

In a recovery situation there is no need to reassemble the operating system, reset configuration settings and restore data.

Restoring a virtual machine from an image-level backup is far faster than initiating bare-metal recovery of a physical server. Virtualisation eliminates the need for one-to-one physical mirrors of a system for Disaster Recovery.

IT has the choice of establishing physical-to-virtual (P2V) and virtual-to-virtual (V2V) failover configurations (locally or remotely) to enable rapid recovery without incurring the expense of purchasing and maintaining identical hardware.

Virtualisation offers flexibility in configuring active scenarios (such as a remote or branch office acting as recovery site for the production site and vice versa), or active-passive (where a corporate-owned or third-party hosting site acts as recovery site, remaining dormant until needed).

Finally, virtualisation delivers flexibility in the form of DR testing. To fully test a Disaster Recovery plan requires disabling the primary data centre and attempting to fail over to the secondary.

A virtualised infrastructure makes it significantly easier to conduct frequent non-disruptive tests in ensuring the DR process is robust and that staff are practiced in executing it consistently and correctly, including during peak operational hours.

With server virtualisation, a greater degree of DR agility can be achieved. IT's ability to respond to service interruptions can be greatly improved, especially with new automation techniques such as those available for VMware virtualisation technology and Microsoft System Center Virtual Machine Manager.

Cloud-based Disaster Recovery

Cloud storage and computing services offer a number of alternatives for cloud-based DR, depending on recovery time and recovery point objectives that a company stipulates.

Cloud computing, along with mobility, social and analytics, accounts for much of the 'hot topic' conversation in IT today. However, when it comes to vendor hype, wild promises and opinions not based in fact, cloud computing appears to stand head and shoulders above the competition.

Because of press coverage that cloud computing receives, many IT Directors have overlooked its usefulness as part of an integrated Disaster Recovery strategy. For IT departments who may not have large budgets to spend in support of Disaster Recovery planning, cloud computing presents an attractive alternative.

This is because the usage-based cost of cloud services is well suited to DR where secondary infrastructure is not used regularly.

Cloud-based DR sites reduce the need for space within data centres, IT infrastructure and IT resources, leading to significant cost reduction. Such sites may also enable businesses without access to large IT budgets in deploying required Disaster Recovery options previously thought to be out of reach.

Many cloud solutions are fixed in terms of services they provide, and may not offer the right level of security or access required by your business. Always ask your provider the following to ensure the basics are covered:

- + How is data securely transferred and stored in your cloud?
- + How are users authenticated?
- + Are passwords the only option or do you as the cloud provider offer two-factor authentication?
- + What regulatory requirements do you adhere to?
- + What bandwidth requirements are needed for the IT team to access data stored in the cloud?
- + If a disaster does occur, do we have the bandwidth and network capacity to redirect all users to the cloud?
- + If we plan to restore from the cloud to on-premises infrastructure, how long will that restore take?

Reliability of any cloud provider, its availability and its ability to serve your users during the course of a disaster are other key considerations to be made.

Defining the DR blueprint with cloud computing

As with traditional DR, there is no single blueprint for cloud-based Disaster Recovery.

Every business is unique in the applications it runs and the relevance of these applications to its organisation and industry sector. Therefore, cloud Disaster Recovery plans should be highly specific for each business.

Triage remains the overarching practice used in defining traditional and cloud-based DR plans. The process begins with identification and prioritisation of applications, services and data, determining for each the amount of downtime that is acceptable before the risk of significant impact on the business.

Priority and required recovery time objectives (RTOs) then determine Disaster Recovery approach.

Identifying critical resources and recovery methods are the most relevant aspects of this process, since IT teams need to ensure that all critical apps and data are covered within the DR blueprint.

To control costs and ensure speedy and focused recovery when the plan requires execution, IT teams should ensure they exclude irrelevant or non-critical applications and data. The more focused your DR plan is will determine how well it may be executed within defined objectives.

With applications, services and data identified and prioritised, and your RTOs defined, the IT team may determine the best and most cost-effective methods of achieving these RTOs, usually on an application-by-application or service-by-service basis.

In rare cases, IT teams have a single DR method for all applications and data; the most likely scenario, however, is having several DR methods protecting clusters of applications and data with similar RTOs.

At Wanstor, we believe a combination of cost and recovery objectives will inevitably drive different levels of Disaster Recovery.

Cloud recovery options

Wanstor business and technology specialists work with you to understand your business requirements, service-level objectives and deployment choices, and match the right mix of delivery models to your business needs.

Managed applications and managed DR

An increasingly popular option is to move both primary production and disaster recovery instances into the cloud and have both handled by an MSP. In doing this, IT teams reap all the benefits of cloud computing, from usage-based cost to elimination of on-premises infrastructure.

The choice of service provider and the process of negotiating appropriate service-level agreements (SLAs) should be top priorities. By handing over control to the service provider, you need to be certain the MSP can deliver uninterrupted service within the defined SLAs for both primary and DR instances.

Backup to and restore from the cloud

Applications and data remain on-premises under this approach, with data backed up into the cloud and restored to on-premises hardware when disaster strikes. In essence, your backup in the cloud becomes a substitute for tape-based off-site backups.

When contemplating cloud-based backup and restore, it is crucial to understand aspects of both backup and restore.

Backing up into the cloud is relatively straightforward, and backup application vendors are extending backup suites with options to back up directly to popular cloud service providers such as Amazon and Microsoft.

The challenging aspect around using cloud-based backups for disaster recovery remains recovery itself. With limited bandwidth and terabytes of data, getting data back on-premises within defined RTOs can prove challenging.

Some cloud backup service providers offer an option to restore data to disks, which are then sent to customers for local on-premises recovery. Another option is a large on-premises cache of recent backups that can be used for local restoration.

Using this approach, data is not restored to on-premises infrastructure, but rather to virtual machines in the cloud. This requires both cloud storage and cloud compute resources, such as Amazon’s Elastic Compute Cloud. The restore may be made either when a disaster is declared, or on a continuous basis.

Pre-staging Disaster Recovery VMs and keeping these updated through scheduled restores is crucial in cases where aggressive RTOs must be met. Some cloud service providers facilitate bringing up cloud virtual machines as part of their DR offerings.

Replication to virtual machines in the cloud

For applications that require aggressive recovery time objectives and recovery point objectives along with application awareness, replication is the option of choice for data movement. Replication to cloud virtual machines can be used in protecting both cloud and on-premises production instances. In other words, replication is suitable for both cloud-VM-to-cloud-VM and on-premises-to-cloud-VM data protection.

Don’t ignore the fundamentals: Blend with new approaches

Cloud computing extends disaster recovery options, can yield significant cost savings and enables DR methods for businesses who lack large IT budgets – it does not change the fundamentals of having to devise a robust disaster recovery plan and testing this rigorously.

	Managed Primary & DR Instances	Cloud-based Backup & Restore	Replication in the Cloud
Instances	SalesForce CRM Email in the cloud	On-premises into the cloud Cloud to cloud	On-premises into the cloud Cloud to cloud
Merits	Fully managed DR 100% usage based Least complex	Only requires cloud storage, virtual machines are optional Usually less complex than replication	Best RTOs and RPOs More likely to support application-consistent delivery
Caution	Service-level agreements define access to production and DR facilities	Less favourable RTOs and RPOs than replication	Higher degree of consistency
Method of Implementation	Accounts Payable	Backup applications and appliances	Replication software Cloud Gateways Cloud Storage such EMC Atmos and Hitachi HCP

Disaster recovery best practice tips

This section outlines the essentials of disaster recovery and business continuity planning for medium-sized businesses.

Disaster recovery is not an easy topic to address. However, by implementing key best practice around disaster recovery, it is possible for IT Managers to know that they will be able to recover from an outage.

Equally, there remain instances where IT Managers lack the knowledge of how to build a disaster recovery strategy.

When scoping a DR plan, businesses need to follow the essentials for disaster recovery planning. The most important and difficult step in disaster recovery planning is to understand how an unplanned outage will affect your business. This step is referred to as a business impact analysis, or BIA.

Without the ability to determine impact of an unplanned outage in a meaningful way, it becomes extremely difficult to determine what type of disaster recovery strategy may be best suited.

At Wanstor, we believe an 'unplanned outage' refers to any unforeseen event that interrupts normal business activity for a period of time, including IT systems failure, fire, power outage or natural disaster.

Depending on the nature of the interruption, this may cause loss of revenue, impact on customer satisfaction, lead to lost opportunities or even result in permanent closure.

Likely outcomes should be determined by identifying critical business activities or functions and extrapolating the outcome of stoppages to these activities or functions – a process where many inexperienced planners fall short.

Quite often, people writing disaster recovery plans skip these important steps at the outset and move directly to articulating solutions, where they feel most comfortable.

Disaster recovery planners should not assume there is to be a workaround or contingency available when highly critical functions go offline.

The intention is to set a recovery time objective (RTO) which covers the period that a process may remain inactive or offline, and a recovery point objective (RPO) which refers to how much data may be lost for critical functions and IT infrastructure to remain functional.

Businesses must determine:

- + A financial value for critical functions, based on financial loss when revenue streams are interrupted. Finance teams should be able to assist with this
- + How critical each function is to the business, based on how each affects revenue streams by using a rating system – for example, one to five, with one the most critical and five the least critical
- + How long a business function may be interrupted before it begins affecting revenue streams
- + How much client or business transactional information may be lost or re-created without seriously affecting the business
- + IT infrastructure and systems upon which business functions depend

The next step is risk assessment, which supplements your impact analysis. The impact of an outage and the anticipated risk that may exist will indicate need for a recovery strategy. Assessing risk is another area where planners may find themselves stuck.

Do not attempt to calculate risk based on the possibility that it could happen, or to calculate annualised loss expectancy. Keep it simple and be realistic about those risks your business might face, including specific threats tied to geographic location.

A risk exists for a business if there is nothing in place to maintain or quickly recover a critical function.

If a system identified as critical is found to have adequate redundancies and protection in place, you may move on to the next system or application.

Once critical functions and their supporting infrastructure have been identified with the impact of an outage quantified using monetary value or a rating system, a recovery strategy can be developed to help prevent or mitigate loss.

It is at this point where IT Managers should consider existing contingencies or redundancies already in place. Critical applications hosted by a service provider under a service-level agreement, probably require little or no recovery strategy.

A recovery strategy is required for applications that support critical functions but lack provisioning in remaining operational.

Specific recovery strategies are determined by a businesses anticipated financial loss if critical functions become unavailable as well as time required to recover necessary applications.

An application with an RTO of within five days may suffice with a tape backup process, but an application required online within an eight-hour business day may require remote data replication and / or standby IT systems at a recovery site.

Outsourcing disaster recovery is also a viable strategy. Companies that cannot afford the cost of developing their own strategy may consider paying for DR availability services, or subscribing to '*DR as a service*'.

The key is remembering that total cost of a recovery strategy should never exceed those losses it is designed to prevent.

The next and final step is to document your recovery strategy and procedure, forming the foundation for a disaster recovery plan.

A word of advice: keep it as simple as possible. Very detailed disaster recovery plans take time to develop and are difficult to maintain.

At a high level, the disaster recovery plan should outline priorities for system recovery, the RTO, recovery procedures, as well as location of data backups and contact information for key recovery personnel.

Testing your plan frequently will help to identify which elements are missing and require inclusion instead of identifying these issues during an actual disaster event. Every time a recovery procedure is tested, gaps are identified along with areas for improvement.

This is how plan maturity is eventually achieved.



Why Wanstor?

An information technology disaster recovery plan should be a *must have* for all businesses, developed in conjunction with the business continuity plan.

Organisations of all sizes create and manage large volumes of electronic information or data, some of this being vital to the survival and continued operation of those businesses.

The impact of data loss or corruption from hardware failure, human error, hacking or malware can be significant – a plan for data backup and the restoration of electronic information is essential.

At Wanstor, we have developed and maintained IT disaster recovery plans for hundreds of businesses across the UK over the past twenty years.

If you need help in shaping a disaster recovery plan that will work for your business should the worst happen, contact us on 0333 123 0360, email us at info@wanstor.com or visit our website at www.wanstor.com for more information.

Request a call back

Wanstor | 124-126 Borough High Street | London | SE1 1LB | 0207 592 7860 | info@wanstor.com

wanstor