



---

## Disaster Recovery – Outsource or In-House?

### A Management White Paper

*Everyone needs a plan*

*Out sourcing DR*

*Priority of services*

*In-house DR*

*Challenges*

*Critical services*

*Second phase of recovery*

*Final phase of recovery*

*The Role of the Third Party DR Supplier*

**By John M Hudson**

Director of T360

*April 2003*

### Everyone needs a plan

This is the initial premise of an earlier white paper<sup>1</sup> by T360. It defines why “Everyone needs a plan”; the assumption is, therefore, implicit.

This white paper looks to examine the options available and evaluate their effectiveness to deliver a real business solution should a disaster occur on the primary site; local and logical threats are covered in<sup>1</sup>.

As a specialist business continuity and disaster recovery company, T360 has regular discussions with a cross section of medium to large organisations about their DR plans. Worldwide events determine where the subject appears on the business agenda, but short-term enthusiasm is rarely sufficient to deliver a cohesive strategy. There is, of course, one consolation; if you know you are without a plan, you at least know where you stand.

### Out sourcing DR

There is a line of thought where the perceived objective of a disaster recovery plan is to wheel in a copy of all the servers, load applications and data: job done! Contracts offering servers to be delivered to a remote site, within say 24 hours, are being considered.

“Without servers, we can’t offer any IT service.” Absolutely true; but what about the rest of it – servers need users, thereby PCs, a network, local back up and recovery. Then there is all the other things essential to users like telephones, desks and chairs and, last but not least, the users themselves; a remote location has logistical considerations.

So a strategy based purely on servers is a point solution, which has little point!

Another aspect common to this approach is “one fit all”, same availability and recovery.

Ask a couple of questions. “How long did it take for the IT department to set up the current servers to offer the service you now experience?” At best weeks, and often it’s been an ongoing process. Is it possible, therefore, to expect a complete change in the IT process that would allow the same group of people to rebuild everything in one or even two weeks? No, it cannot be achieved. So why have all the servers on a rapid availability contract at vast expense?

Another area of concern arises when the plan in place cannot or has not been tested and on the day it is called into action it fails or falls short. Today, greater value is placed on the data recovery process using commercially available tools, which makes the solutions more robust, but without testing there can be no guarantees.

It is for this reason that plans incorporate periodic DR testing to demonstrate to the business that services can be recovered - don’t knock it, the concept is a very valid one. However, is the test valid or representative in a real scenario? The truth is that they are often inconclusive.

Through an outsourcing contract the test is carried out on a third party site for which a forward booking has been made allowing IT staff to prepare and, during the test, go back to HQ for forgotten CDs or documentation. The servers are ready but devoid of application; often the first task is to build a media server to enable data to be recovered. Data is then poured onto the servers, and the test timed, though still without an application present. The charges incurred for running these tests limit the time available, rarely allowing for operability testing or to have real users connected to them. Can this artificial environment be considered a real DR test? What has really been achieved? The data can be read from tape, proving the recovery software and library both work and the manufacturers of both those products would confidently endorse their respective functionality. A complete re build with serviceability testing is the only conclusive way to ensure the integrity of the resultant end user experience.

For larger organisations the DR testing is segmented into business verticals, so a significant element of the IT team is recovering a small part of the business’ IT infrastructure on a regular basis. In a real disaster, the same team would be stretched across the whole business and asked to deliver a lot more.

<sup>1</sup>Disaster recovery by replication; another white paper by T360.

## Disaster Recovery – Outsource or In-house?

Outsourcing DR services has a role, as discussed below, but there are three areas of concern:

1. A blanket contract will not deliver against the business objectives; there is payment for something that cannot be used.
2. DR testing on their site achieves little but costs in resource that is not available to the business; 100 man-days costs more than £50,000 even at in-house rate.
3. If the balloon really goes up, can they deliver against many contracts simultaneously for example, 9/11?

The conclusion must be that this is a cost without a return.

Whilst IT vulnerability is not acceptable, neither is wasteful IT expenditure. Business demands more for less from IT. Therefore, the key is to put in place a strategy which delivers IT service to the business, with minimum impact and cost, that can be deployed to agreed operational objectives, and with the available resource.

Therefore, the first challenge of the plan is to give each service and application a priority.

### Priority of services

Whilst an integral part of the business, IT cannot be expected to make these decisions in isolation. These are business decisions with fiscal implications and as such require ratification at the highest level; sign off at Board level is essential.

Every application and IT Service has a value to the business, a service business value (SBV) quotient, derived from a combination of factors:

- Business revenue delivered
- Business impact
- Application complexity

Customer facing services generate direct revenue, B2B or B2C, so can customers place orders without this service?

When would the absence of this service hurt the business? This is less clear-cut.

A billing system generating invoices, the lifeblood of the business cash flow, at first glance is critical but closer inspection shows this is not the case. Legally, invoices can be issued later and not impact the due date, many are “reissued” as they were “lost in the post” anyway. The important thing is that there is no impact on the business to trade or on customers.

Customer facing services, however, do have a huge potential to adversely impact the business, even if they don’t generate revenue directly. The absence of service will impact the quality of customer experience<sup>2</sup>, resulting in churn; this in itself has large fiscal implications.

A complex application may not fall into the first two categories but, by definition, requires more resource to maintain and take longer to recover than a simple one. In order to meet its operational objectives, the recovery process must be started earlier, artificially raising its SBV.

In summary, when will the absence of each service hurt the business fiscally or prevent it trading?

The loss of internal systems is rarely the threat first thought, whilst customer-facing services would yield a quicker impact, though interestingly businesses don’t always recognise this either.

Defining the priority of services is an iterative process, as the initial list of “must haves” is too excessive on cost or resource ground, third and fourth passes are common. Be absolutely ruthless.

<sup>2</sup>Service Assurance; Another white paper by T360.

## Disaster Recovery – Outsource or In-house?

The process of analysing each service against these criteria often leads to strife.

Departments inflate their own IT value to the business. This posturing is in direct conflict with the business objective that is to minimise the number of IT services deemed crucial with a high SBV.

Service level agreement (availability and back to service / recovery), and quality of service, (perceived end user responsiveness), are derived from and delivered against the SBV. This in turn defines the infrastructure required to deliver the service from all aspects of operation.

### In house DR

The economic climate does not allow for massive additional investment, let alone the provision of a secondary DR site fully equipped with on-line replication – a common assumption is that this is a non-starter. Well, maybe not.

The cost implications are staggering *if* the aim is to duplicate the primary site. The bandwidth required to transfer every single data transaction from every user on the primary would be frightening in itself!

Therefore, based on the value to the business of each application as discussed above, be strategic.

The IT team is finite, which can recover applications and services in a given time, against a plan defined by the business. Herein lies the key. Replicating everything is as meaningless as having all the servers delivered to the remote site at the same time. Make the strategy a phased approach and then make what has been invested work for the business. Design a deterministic solution; give the business full control of both the operations and costs. Use the investment daily; don't merely pay it out as an insurance premium; it is dead money.

### Challenges

The business has defined its service priorities, and has a clear understanding how long it takes for the IT department to build each of them from scratch to be fully functional and tested. From tape back ups, this is rarely hours, and often days.

Two issues will arise. The first is a skills shortage; a member of the technical team cannot be in two places at once, and larger applications require several people at different times to complete the process. Some additional training and skills transfer can address this.

The second is more crucial. If the back to service agreement (BSA) for a high SBV is four hours, by recovering from tape, this objective will not be met unless the data set is small.

This places another impasse when recovering on a third party site. The first task is to build a recovery engine, installing Veritas, Legato, or other back up software tools, with a suitable tape library, network it in some way to the new target servers, SAN, LAN, or NAS and ensure it is all operable. Result; over half a day is lost *before* the recovery process starts. With a 48-hour BSA, even a modest Oracle application will not be completed within the operational objective, and nothing can be recovered in a 4-hour BSA.

### The Critical Services

This list of critical services must be as small as possible. When it is defined, ask the question again and reduce the list further. This is the expensive element of business continuance.

These applications must have significant value to the business, and as such be viewed as high availability (HA) services. To achieve a rapid return to service, the data must already be available. For this replication is required, preferably to a secondary site you own or rent.

The strategic deployment of loosely clustered servers, utilising either synchronous replication over fibre or asynchronous over IP to transfer the changes will provide resilience. By making this bi-directional, the second resource can either share the load or deliver another HA service back the other way. If the data sets are minimalised<sup>3</sup>, then the connection between the sites will cost less.

<sup>3</sup>Cotton wool around data back ups; Another white paper by T360

## Disaster Recovery – Outsource or In-house?

To reduce the investment cost, examine if the quality or functionality of the service can be reduced. A subset of the application or data, the secondary servers can then be a lower specification or shared by more than one application. Provision needs to be made for a minimum infrastructure to support it, a small user network, a back up device, after all this could be home for a little while.

A benefit of HA is derived from reduced downtime associated with system and software upgrades, where an equally challenging BSA must be addressed several times a year.

There are operational benefits of owning the servers:

- They are set up ready with the application configured
- Their status is known, operating system, patch level
- They can be tested regularly and easily
- They can be used to deliver value to the business
- The time to recover is minutes not hours or days
- They are built outside the DR plan, reducing risk and business impact to the most critical services

The business benefits based on reduced costs and ROI are clearly visible:

- + Operational objectives (SLA and BSA) can be met and even bettered, reduced loss of revenue and business impact
- + It can be designed so that no additional servers are required
- + Any further investment made is in your business not a third party's
- + The business has control without dependency on a third party
- + Reduced cost to a third party more than off sets additional servers
- + Resilience through change or upgrade
- + Changes can be carried out in office hours

The costs associated with the link will be reduced if a minimalist view of replicating only essential data is taken. Surplus bandwidth in the Telecommunications industry at this time makes this process even more viable.

Replication has always been viewed as complex and expensive, by being strategic, this is no longer the case.

### The Second Phase of Recovery

Having restored the critical services; the 10% of systems that impacts on 90% of the revenue: what about the rest? A strategic approach is required. The remaining services will have to be built from scratch almost certainly from tape back ups.

Some of the IT team must be engaged in setting up user accounts, systems management and initiating a back up process on the new primary site, they therefore cannot be included in the planning for the recovery of the remaining service.

The objective is to determine the next tranche of highest priority services that can be recovered in a two-week period with the available resource. Two weeks is a typical delivery time for servers from a manufacturer; this generalisation requires tuning for the businesses specific circumstances.

At this point there is a need for the out source DR supplier to deliver appropriate servers to site, but this is but a subset of those initially envisaged and based only on those required for this short phase of the DR plan.

The recovery engine is already in place and the process can commence immediately. Again, as this is under the businesses own control, it is easier to implement a number of automated mechanisms to speed this process. These are outside the scope of this document, but suffice to say every business is different; there are no generic off the shelf solutions.

### The Final Phase of Recovery

And then there are the rest, which might equate to a significant number. Two weeks in, the full extent of the disaster has been assessed, the duration to the transitional location will be clearer, and a medium term plan in place; which could include another move of data centre!

By agreement with the business, these services have been designated as a lower business priority, by definition; the business can function without these services for more than two weeks as part of the BSA analysis. The equipment for this phase can be purchased, delivered in readiness for restoration by the IT team, and the costs borne by the company insurance policy. The key here is that the business has bought time, has all its options and, most importantly, kept money in the bank.

### The Role of the Third Party DR Supplier

Providing servers and applications is but part of the disaster recovery strategy, and at this point specialist third party outsourcing agencies can provide additional valuable help.

- Desktops systems
- Telecommunications
- Desks and chairs
- Network
- Servers for phase two

If the organisation operates a twin active site policy the loss of a site will still offer a skeleton facility in place from day one. Deploying many generic desktops can be easily automated and, having them set up identically has many operational advantages. Arranging for the re-routing of telephone lines, transparent to the customer, requires the facility available only to a telecommunications operator, or its agent, and cannot be addressed in house.

Two areas common in existing outsourced DR contracts can cause significant problems and should be addressed before commitment. The business will change, dramatically in IT terms over a 3 or 5-year contract period, through growth or obsolescence. The DR contract must have the flexibility to embrace these changes, and be reviewed quarterly. Secondly, who establishes that it is fit for purpose? Many DR suppliers use “tick box” techniques, any omissions are the responsibility of you, the customer. This is a questionable approach; after all they purport to be the experts and as such should sign up to the responsibility of due diligence.

### T360

T360 is an innovative solutions provider that assists organisations to transition IT from being a necessary component to a strategic collaborator within the business.

#### **‘Aligning IT to the business is a Process not a Product’**

Based upon best of breed ITIL solutions, T360 enable customers to achieve an end to end visualisation of their business services delivered by IT through the areas of:-

- Business Service Management
- Business Activity Monitoring
- Customer Experience Management

T360’s philosophy is simple; we listen, understand and deliver. Our customers trust us because we provide innovation, expertise and commitment.