Typical disaster recovery tests – a waste of time?

How would your disaster recovery provider cope in the event of a real disaster and are your IT managers up to scratch? Dave Claridge proposes a change of strategy

Disaster recovery tests are used to convince businesses that the IT disaster recovery plan would work in the event of a real disaster. But in reality many of these plans would never work with the consequence of a significant risk to business continuity. This article explains why and proposes a new strategy for disaster testing. We will also look at a case study of an insurance company that used this strategy for the first time and achieved a recognised formal accreditation for its disaster plan.

What's wrong with disaster recovery?

The majority of companies with large computer systems that are critical to the business have some form of IT disaster recovery plan and process in place. However, few of these would actually work in a real disaster because:

- processes do not exist for ensuring that the disaster infrastructure and plan are up to date;
- recovery plans are incomplete; and
- there is no link between the IT disaster recovery plan and the business continuity plan.

The IT department and the disaster recovery supplier have just one aim in mind – "to make sure the disaster recovery test works, no matter what".

This leads to detailed pre-planning of tests such as booking transport for the shipment of tape media, making last minute changes to the IT recovery environment, arranging hotels, people and so on.

But what is actually being tested? Simply, nothing of value other than the IT department's ability to recover the systems given sufficient time and warning. Of course, in a real disaster, neither would be available.

What is needed is a new strategy where disaster testing simulates a real disaster as far as possible. This strategy also needs rigorous processes in place for continuously maintaining the currency of the recovery infrastructure.

Insurance company case study

This article looks at an insurance company which experienced significant problems with disaster recovery testing. The insurance company had a complex environment consisting of a leading edge high-speed switched LAN (local area network) running a workflow and image system linked to a mainframe computer. This ran core business systems on CICS and IDMS, with a UNIX system providing decision support. The services were provided to a number of branch offices around the UK via a WAN (wide area network).

Disaster recovery services were handled by a specialist disaster recovery provider. Business continuity plans were not yet in place and IT recovery plans were out of date and incomplete. We will focus on the recovery of the mainframe services, the network. Image and workflow recovery was not part of the brief at the time.

Business requirements

Loss of the computer systems would leave the insurance company unable to operate. The company needed the critical CICS service to be restored within 15 hours to continue operating. The original disaster recovery plan could only achieve 48 hours – and even then disaster back up for the network was not available, so connection to the recovered system would be delayed by the need to buy and configure network equipment.

What this meant was that, although the insurance company had invested in disaster recovery, the investment was mis-directed and effectively wasted. ■

The first disaster recovery test

Shortly after starting work with the insurance firm to implement the new strategy, we observed a disaster recovery test in progress. This was the first to be carried out for two years. But since the last one, the firm had migrated to a newer version of its operating system and made a number of hardware changes.

During the test we observed a number of significant problems, including site access problems, incorrect hardware configurations, incomplete documentation and software glitches, which rendered the test of no real value other than showing that the process simply would not work in a real emergency.

The postmortem

We spent time reviewing the results of the test with the insurance company and proposed a number of recommendations, as follows:

- evaluate other disaster recovery providers in terms of facilities and support;
- implement a disaster recovery problem management process so all problems can be traced back to root cause and prevented from re-occurring;
- produce new comprehensive operations documentation on how to start up the systems;
- organise support resources in a separate area away from the operations bridge and formalise communication

of problems (such as telephone and problem report);
- modify the change management process and include special provisions for disaster recovery;
- automate the start up of the system as much as possible so as to simulate normal production running and avoid intervention;
- implement a continuous formal paper-based process for communication with the disaster recovery provider triggered by any relevant IT changes; and
- start work on a comprehensive disaster recovery plan.

Follow up activities

The disaster recovery vendor – Our view, shared by the insurance company, was that the current provider could not support the firm's requirements for disaster recovery. A number of providers prefer not to get involved at all with their customers' tests. But we believe this is not the best approach. Its criticality to any business demands a partnership between the company and its disaster recovery provider.

Bids from the providers vary enormously. Changes in technology and the reduction in hardware costs mean that new disaster recovery contracts are now much cheaper, and the business was won by a provider which acknowledged this in the quote.

Problem management processes – The disaster recovery test was thoroughly documented in terms of a description of all problems that occurred, including those fixed during the test, no matter how trivial.

Each problem was assigned to an owner with an open status and with a target date for resolution. The dates set were aggressive and not related to the date of the next test. This is essential because while problems remain unresolved the business is exposed to a major risk should a real disaster occur.

Resolving problems is not just about fixing the problem – a process must also be in place for preventing the problem from recurring.

Recovery processes – During the first test the systems programmers closely watched the operators and at the first sign of a problem took over from them. We also noted that one of the operations staff had attended all the tests and therefore relied on his knowledge of what had happened previously. We also found that the documentation was out of date and difficult to follow.

Resolving disaster recovery problems is not just about fixing the problem – a process must also be in place for preventing the problem from recurring

There are two lessons here. First, the operators must be responsible for starting up the systems and they should instigate any need for the involvement of systems programmers. Second, it is important to rotate a number of different operators through tests, and that the operators follow step-by-step instructions. These were built into the next test.

Change management processes – The problems from the last test showed that one of the key issues was to ensure that any changes made to the production infrastructure are also reflected in the disaster recovery infrastructure. To provide this the change control procedures were adapted. In particular, the change control form was modified to include a special field relating to disaster recovery so that if a change had an impact, the corresponding change had to accompany the change requested.

Communication with the disaster recovery vendor – A key testing issue is that vendors are very keen on meetings immediately prior to the test itself. This will do nothing to help prepare for a real disaster. Preparation for real disasters must be ongoing.

A formal process of communication with the vendor was established, and this new process checked for any changes to the IT infrastructure or the recovery plan. If these affected recovery then an updated set of the corresponding documentation would be sent to the vendor without delay (the same day). The vendor would then be responsible for adjusting the details of the insurance companies configuration to keep it synchronised. The vendor would confirm that he had reflected the changes as a check and balance.

Streamlining the recovery process – The insurance company's IT infrastructure was not up to date in terms of magnetic tape technology. This meant that a larger number of tapes were needed for the recovery back ups. By using tape hardware compression, the company was able to reduce the number of tapes and also reduce the back up and recovery time.

The second disaster recovery test

This test was far more successful. There were fewer problems, but there were still some issues to be resolved.

Change of disaster recovery vendor

Shortly after the second test, notice was served on the recovery provider and a move made to an alternative provider. The next test would be with the new firm, giving us time to resolve the issues surrounding the recovery time.

Reducing the recovery time

The new company used a tape storage provider that was some hours away by road further away from the original provider.

This is a common situation with a lot of IT installations, and is often based on an assumption that the provider will not be ready for some hours after being alerted to the problem. But this does not have to be the case. The provider may have the capability to set up a configuration very quickly. There is also the possibility that the transportation of tapes will hit serious problems, for example, due to an accident, a road closure or bad weather. Therefore, there is a significant advantage in storing tapes with your recovery provider. This also allows them to recover the systems on your behalf, say, if you have lost a number of your key staff as a result of a disaster.

Preparation for the next test

The insurance firm's plan was incomplete and out of date, so we began work on creating a new plan covering in detail the 'who does what' processes, which were laid out in a check list format required to carry out

the recovery. It was also important that holders of this plan kept it somewhere readily accessible and not at their place of work.

The third disaster recovery test

The third test was very successful, with only a few, mainly trivial problems. We decided with the provider that it would kick off the initial recovery of the data for all future tests. This was possible because the provider also held the recovery back up tapes. The system instructions were also written clearly so that someone who wasn't aware of the specifics of the insurance company's systems could still recover them.

Managing the disaster recovery provider

One of the things we were keen to do was to get the provider working in partnership with the insurance company, so a continuous improvement of the disaster testing process could be achieved. We focused on several items with the provider associated with the testing process. Effectively, what we were doing was drawing up the processes for best practice recovery testing and getting the insurance company and the provider working together to achieve this.

The role of internal audit

Internal auditors often use a process of interviews to carry out an audit of disaster recovery. However, we felt we could improve upon the effectiveness of the audit by inviting the auditors to observe the next test.

We had produced a set of score cards for recording the results. These score cards covered the key activities involved – 16 items. Each item was documented with examples of how to judge how well it had been achieved.

'No warning' disaster recovery test

The next test occurred without any warning to the IT staff. Only two senior IT managers would be aware of it. Internal audit used the results of the test to cover their audit and used the score cards to assist with the process. A key part of the test would be to monitor how effective the new plan would be.

We planned that the IT technical support manager would call a meeting about IT strategy with all of his managers. The meeting was to occur at 8am at an offsite location. During the meeting, the technical support manager would declare a simulated disaster and ask his managers to recover the systems by 7pm that evening. At this point he would walk out of the meeting and be unavailable all day.

The first 'no warning' disaster recovery test

We learned a lot from this test, particularly about the overall logistics of the recovery and what could go wrong. In particular, it was observed that the managers took one-and-a-half hours before they decided to invoke disaster recovery services.

Essentially, they didn't attempt to get access to the recovery plan. Instead, they tried to remember what occurred in the previous tests. In a real disaster this would have caused the loss of a crucial amount of time.

Interestingly, there was significant feedback from the emergency management team regarding the disaster recovery plan. In attempting to use it in a simulated disaster situation, they learnt a great deal about ways to improve it. In particular, issues came up about the availability of car hire companies, credit cards, mobile phones and similar logistical items. Updates were made to the plan in line with the feedback.

One innovative idea was the concept of a credit card-sized recovery information card. This would hold details of the location of the plan, the invocation process and contact numbers for the emergency management team. This card was produced and used successfully on the next test.

The second 'no warning' disaster recovery test

For the next test, internal audit would make a call to one of the emergency management team at 6.30am. The instruction would be that a disaster simulation was in progress and that the systems need to be recovered by 7pm. The test went extremely well, with all the data recovered on time and intact.

Issues with 'no warning' disaster recovery tests

One of the key reservations cited by a number of companies for staying with planned disaster testing is that IT or business management will not support the concept of 'no warning' tests. The reason for this is that any deflection of day-to-day resource to support a 'no warning' test is seen as having a potentially adverse impact on the service being provided to the business.

We believe that such an impact can be justified to the business on the basis of the importance of disaster recovery and the need to get value for money from the significant financial investment made in it.

Strategies for world class disaster recovery

There is still a role for the planned disaster recovery test under certain circumstances, for example, a major IT change, such as a new operating system or complex application. However, there is no need to involve the recovery team for this, and the provider can perform the tests, given appropriate access.

Other tests should be called without warning and should occur randomly. An independent and suitably authorised party such as internal audit should organise these. They would not know about the tests and only one person in disaster recovery should be aware.

Disaster recovery must be considered at all times, particularly when cost benefit cases are being carried out for new applications.

An effective plan

What this article has shown is that a rethink of disaster recovery testing strategies is needed. We advocate that the concept of planned testing is totally changed to that of random 'no warning' tests where what is being tested is what may well be experienced in the event of a real disaster.

This is as opposed to what typically gets tested, which is how well the IT department can perform if it gets a chance to plan the test and it is given sufficient warning.

If the strategies for world class recovery discussed in here are followed, your plan will be truly effective. Should a real disaster occur, your business will be protected. In today's somewhat uncertain world where, eg terrorism, is prevalent this is essential. ■

Dave Claridge is Principal Consultant at KPMG Consulting.

www.kpmg.com