

# LESSONS FROM THE HOT SEAT: An Insider's View of Disaster Recovery Testing

Daymark Solutions, Inc.

## On the Hot Seat: A Real-Life Tale

*It's 3 a.m. and here I sit watching the bits increase and the time counters advance. This is about as productive as watching paint dry...*

*I'm in a hotsite disaster recovery center watching a set of restores run for a client's disaster recovery test. It's not fun, it isn't going well, and no one is happy. The only saving grace is that we were able to convince the client to split up into shifts so that some people could catch some shut-eye while the rest of us charged on into the night. We finally got the first round of restores running about two hours ago but we were only able to get half of the hardware functioning.*

*We are halfway through the client's allotted time at the DR center. We're already halfway behind, and now we can only hope to run at a maximum of half speed. It looks like it is going to be another long and sleepless night.*

*I ask myself yet again: "What is this fire drill really accomplishing?"*

*I flew into Philly late last night and met with the group of computer administrators at the hotel bar. I know most of them fairly well, as we had worked together on other projects in the past. All of them work for a large financial company that has a high-level support contract with my company. I'm what they call a "fly-to-site" support technician. I've been doing this type of work for a while now. I fill the role that's advertised as "business critical support." Some people refer to it as smoke jumping and a lot of the guys who do this type of work like to call themselves cleaners. In between flying around the country to put out various computer issues, I also assist with disaster recovery tests like this one, which has brought me and this team of administrators to Philadelphia. Twice a year we come here and run a mock disaster to test how quickly we can bring the customer's data back online.*

*We started at 8:30 in the morning and hit the first snag by 9:30.*

*During the DR container inventory, it was discovered that the latest operating system patches were missing. It wasn't a big deal, though, since we were able to download them from the web and most everything else seemed to be intact. Fortunately for the backup team the network guys weren't ready yet, so we made a quick decision and sent a team member back to the hotel to download the service packs to his laptop in the hope of keeping things moving at a reasonable pace.*

*By 9:45 a.m. we were loading the backup utility on the "pre-configured" backup server provided by the hotsite. Once the server was up we hit the first show stopper: the operating system wasn't seeing any of the tape drives.*

*At 11:30 a.m. it became clear we were falling behind schedule. To make matters worse, the network team encountered a hardware problem which would mean no network connectivity for another hour or so. When the troubleshooting began on the tape drives, we encountered yet another issue: the tape library we had been assigned was in a secure section of the hotsite facility that we did not have access to. The delay continued as we contacted our liaison, who in turn contacted the DR facility technician to verify the tape library. After an hour and a simple reboot of the tape library, we had half of the drives accessible from the operating system.*



*With the amount of data contained in the application backups, having half of the allotted tape drives wouldn't cut it. There just wasn't going to be enough time to get all of the data off the tapes. We decided to divide and conquer, with one group focusing on the tape drives while another rebuilt the backup catalog in preparation for data restores. More time was wasted when we couldn't reboot the backup server to address the drive issue while the catalog was being built.*

*Eventually the network came online and most of the operating system groups had at least one server ready for restore. Unfortunately, the backup team couldn't get the restores to work. Suddenly we were the bottleneck holding everyone else up. Talk about pressure. After a couple of reboots and some adjustments to the HBA configuration file, still only half of the drives were accessible from the operating system. While one team was attempting to address the drive issues, the other team successfully brought the catalog back online. The first set of test restores, however, did not work.*

*After about seven hours of hard work, there was no data flowing, no restores active, and all of our benchmarks had consistently been missed. We were never able to determine why all of the tape drives were not available. We spent the next twelve hours or so in pure firefighting mode. I remember thinking during a brief soda break around 7 p.m. that this was shaping up to be another all-nighter.*

*And now at 3 a.m., watching the paint dry, I curse myself for being right.*

I've been going to hot sites with clients to do disaster recovery testing since 2000. I was also the lead technician on the Primary VERITAS "swat team" following the September 11th tragedy. I have spent most of my professional career working with data protection and recovery solutions.

The purpose of this paper is to share what I've learned during this time and to offer insights into the importance of effective disaster recovery testing and preparedness. I stress "effective" because perhaps the biggest lesson I've learned is that most disaster recovery testing, while useful in uncovering certain problems, does not accurately evaluate a company's disaster preparedness.

## **The Data Recovery Lessons of September 11th**

Following the events of September 11th, many of the major news outlets were reporting that no significant amount of data was lost. Many of the reports stated that most of the companies in the World Trade Center had adequate backup and disaster recovery solutions in place and were able to recover all of the information. Unfortunately, that was not exactly the case. A more accurate statement is that no data that was physically available after the tragedy was unable to be restored. I can tell you one thing for sure – no two data protection strategies were alike for the companies operating in the World Trade Center prior to the attack.

One company was backing up its NYC data center to New Jersey and New Jersey to NYC. Another was mirroring all servers located at the Twin Towers to NJ and had been since the first terrorist attack on the Twin Towers in 1998. Other companies had small islands of replication for protecting their critical servers only. Most of the mid-sized to smaller companies were doing tape vaulting and shipping tapes offsite to a secure location.



Of the companies that were only using a tape vaulting strategy, all experienced some level of data loss. The amount of data lost by each company was directly related to the offsite tape rotation schedule they had in place.

---

**LESSON LEARNED** If the physical media is unavailable there is no way to restore the data.

One firm, for instance, had an offsite data pickup scheduled for Monday, September 10th, but it was missed due to a mechanical problem with the truck. The last time that firm had sent data offsite was the morning of Friday, September 7th – prior to their weekend full backups.

---

**LESSON LEARNED** The offsite tape rotation schedule is directly related to the RPO that is achievable following a disaster.

This same company handed me a box of 800 DLT tapes with no catalog backup to index them. Without a catalog backup the only solution is to manually import each tape into the backup application and rebuild the catalog. If the backup application in use doesn't support re-indexing the catalog, you are out of luck. Re-indexing the backup catalog can take an inordinate amount of time – time that is simply not available during a recovery.

---

**LESSON LEARNED** Include the backup application in the disaster recovery plan.

One admin at another company knew exactly how important the backup catalog information was and took the time prior to evacuating the building to un-cable the master server and carry it out with him. Unfortunately, they were still only able to recover data from the tape media that was physically available after the tragedy.

---

**LESSON LEARNED** Even with the backup server fully intact and the catalog in place, you can still only restore data from the media that is physically available after the disaster.


The location strategies on where to do the recovery also ran the spectrum. One company had a DR location under wraps and was up and running when the markets reopened on Monday. Many of the large financial companies had alternate data centers they could use, as well as good recovery strategies. Most of the midrange companies had contracts with a hot site data recovery center similar to the one mentioned in the introduction to this paper.

When I went with one company to a hot site in northern New Jersey, it was standing room only. I had to flip a trash can upside down so I could have a place to sit while I shared one modem line with eight other people. While some of the smaller companies at the hot site had a fair amount of success with their recoveries, all of the larger firms ended up leaving rather quickly after determining the facility did not have the resources they needed in order to be successful.

---

**LESSON LEARNED** In a wide area disaster, resources can and will be overcommitted.

One of the most difficult lessons I learned during this time was the long-term impact to the business of losing key people in the organization. We were restoring a set of Oracle databases, a task that was particularly challenging due to the way each database had been backed up. As difficult as the restore was, the hard part of the lesson was being unable to figure out how each database was backed up. Eventually we were able to rollback until we found a good backup to restore from. After the restore was complete and the database was back online, I learned that the entire



team responsible for the database had died in the attack. As a result, no one knew what the data was, what it was for, or why it was being tabulated. Seeing firsthand how the loss of key people affected this company has never left my mind.

**LESSON LEARNED** Document everything. You do not know who will survive a major disaster.

## The Fallacy of the Disaster Recovery Test

*So... it's still 3 a.m. and here I sit. The paint is not yet dry and I'm trying to figure out what the real value of this disaster recovery test is. I have spent most of the day working through configuration issues and hardware problems. I guess there's some training value in this type of firefighting, but is the cost of six to eight people traveling out of the area for two to three days to work long hours, not to mention the direct cost of the contract, worth it? I wonder just exactly what it is that we are trying to test here. Are we validating the Recovery Time Objective (RTO) for the data or hoping to confirm the Recovery Point Objective (RPO)? Which is the most relevant?*

*I wonder how much prep time went into preparing the DR container for this week's test. I overheard one of the administrators talking about how many times they had to rerun a backup last week before it was successful. Would they have done that if the test wasn't this week? Was that one of the critical applications he was responsible for? What if this was a real disaster that happened earlier in the week – would they have lost any data if that backup wasn't complete? I also happen to know they did a special tape offsite for the test directly to the hot site. I wonder if the data being stored at the offsite tape storage facility is identical to the data we're using here for the test.*

These are just a few of the questions that made me realize that too many companies are not really testing the disaster recovery preparedness of their businesses. The testing being done does not measure the RPO of any realistic scenario. Many of the tests I participated in had no criteria to measure how close to the actual point of the disaster the data being used was. All of the benchmarks we missed were focused on the scramble to rebuild the backup application, not to confirm the validity of the data.

The original goal of my client's test was to prove that their top five critical applications were capable of being restored within a predefined Service Level Agreement (SLA) of 24 hours. The SLA was exactly the same for all five applications due to the amount of time the hot site had available for testing. The hot site provided a shared utility for testing that was constantly being reconfigured for many customers. Most of the delays we experienced were due to hardware problems or configuration issues. It ended up costing the client a lot of money to prove that the external environment was the leading cause of delay in the testing.

Needless to say, we didn't accomplish our goal, but at least the client could take advantage of a reduction in their insurance rates by doing a scheduled DR with a qualified third party. Additionally, they found some immediate things they needed to fix – even an unsuccessful test has its advantages. This client was testing to the application level, which does have the value of checking the recoverability of the entire business path to the end user. Just testing to see if the backup server can be rebuilt is more a test of the technical competency of the staff than the disaster preparedness of the company.



## How to Derive Real Value from Disaster Recovery Testing

Perhaps the next time this company can be convinced to do an impromptu test, which is truly the only way to determine just how prepared an organization is for unexpected downtime. Disasters rarely give warning and the most tragic are almost always the most sudden. When creating a disaster recovery plan, the best advice is to “hope for the best and plan for the worst.”

Furthermore, your most talented people should not be involved in every test. These are your frontrunners, the people everyone leans on the most. What if, in a true disaster, they are unavailable? How would your staff cope? Most DR experts strongly recommend that companies rotate their staff so that all members participate in recovery testing to ensure that they all have this critically important experience.

Where would you score in a true **test** of your disaster preparedness? How would you measure it? A true test of disaster preparedness is unannounced; most tests I've participated in have given the team at least a week, if not months, to prepare. An impromptu test would offer more realistic insights into what data is being safely stored at an offsite facility and how well you can recover from a potential disaster.

### Recommendations for Effective Disaster Preparedness Testing

In its many years of architecting, implementing, and optimizing backup/recovery solutions for businesses, Daymark Solutions has developed several tests which have proven to be highly effective in measuring a company's ability to ensure business continuity in the event of a disaster.

#### Onsite Recovery Testing

Daymark provides a preconfigured backup application with tape devices, application hosts, and storage for the restores. Experts from Daymark conduct a test focused on a specific application or subset of applications chosen by the customer. The goals are twofold:

- 1) to train your staff on how to recover the selected application from a bare metal environment; and
- 2) to certify that the RPO and the RTO of the application(s) selected for testing meet any predefined SLAs.

The value of this type of test is that it removes the backup application and associated hardware as factors and offers a true view of your disaster preparedness.

#### Mock Disaster Testing

As with Onsite Recovery Testing, Daymark provides a preconfigured backup application with tape devices, application hosts, and storage for the restores. The client chooses the application or subset of applications to be tested.

What makes this test different – and significantly more valuable – is that it is unannounced to the staff, a true mock disaster situation that can be done at your location or at Daymark's state-of-the-art testing lab.

For the most conclusive results, Daymark recommends a two-phased test that combines onsite testing with mock disaster testing. In Phase 1, Daymark conducts the Onsite Recovery Testing with your staff as a structured training session. In Phase 2, Daymark conducts an unannounced test during which our experts evaluate the staff response and environment for disaster preparedness. It also measures the results against the SLAs of the selected applications.



This manner of testing can also be leveraged for compliance standards such as Sarbanes-Oxley (SOX). A graded test run by an independent third party can go a long way during a compliance audit. By its very nature, an unannounced test helps to ensure data protection and recoverability and can be seen as a “stitch in time that will save nine” when it comes down to a SOX audit. Completion of an unscheduled DR test to the application level satisfies specific SOX audit requirements such as how close to the time of a given disaster the current data set is being used for restore and if there are the required amounts of recoverable datasets for a given application available in the event of a disaster.

The application used for the test, be it Oracle, SAP, Email, SharePoint, Apache or another technology, will have a documented, tested backup and recovery plan that has been verified to be not only capable of, but also in use by, your staff to recover the financial application requiring the SOX compliance within both a predetermined RPO and RTO.

## Summary

DR testing is supposed to give a company confidence that its DR strategy will perform as expected to ensure business continuity. In reality, however, most testing fails to accurately reflect the DR preparedness of the company. At most, it may measure the technical competency of the staff, or uncover some unexpected problems that must be fixed.

To truly measure your ability to recover, you should conduct impromptu tests that evaluate the response of your infrastructure and staff. Since the goal is to simulate a realistic DR scenario, any advanced notice or preparation should be discouraged. Score the response of your staff and your infrastructure, and measure it against specific SLA objectives. Engaging a third-party expert like Daymark Solutions can often yield the most accurate, and unbiased, analysis.

## About Daymark

**Daymark Solutions, Inc. is an experienced technology solutions provider focused on architecting, providing and implementing effective server, storage, network and security solutions for businesses in the Northeastern United States. Our mission is to help you use technology effectively to solve the business challenges you face today while making sure your systems are agile enough to adapt to future requirements. For more information, call 781-359-3000.**