

# Disaster Recovery in the Virtual World

---

Having worked in IT for over a decade both as an internal systems engineer and at a solutions provider as a technical consultant designing and implementing solutions for clients, I have been involved in numerous disaster recovery projects. The recovery of IT systems has always played a major part in the overall planning of businesses Disaster Recovery strategies. Having worked in this field for a long time I am all too aware of the requirements and potential complications of recovering failed systems or a lost site in a physical IT environment.

In the last 3-5 years I have seen at firsthand the changes brought about by virtualisation technology in mainstream businesses, I have seen how companies are benefiting daily from technologies such as snapshots, vMotion, Storage vMotion and easily altered resource allocation to increase system uptime, reduce out of hours working and to increase response times to user and business requests.

However I would say that one of the most compelling areas where we are seeing virtualisation technology truly delivering cost effective benefits is by infinitely reducing the burden on system administrators and preventing wasteful purchasing on hardware and licenses is in disaster recovery.

## *Historically*

For years, disaster recovery planning has entailed obtaining significant amounts of server hardware, either by repurposing servers that have been replaced, or by purchasing new servers for the sole role of sitting idle as a passive node ready for the recovery of the primary system.

Typically the server protection model has relied on a 1:1 correlation between a live business critical server and a dedicated protection server; software solutions such as Microsoft Windows Clustering, software replication or, a combination of data copy scripts, service start-up scripts and a lot of manual interaction have all been used to manage data replication and application services failover and failback.

Many of these solutions involve the alteration of the live business application environment, i.e. Microsoft clustering requires an Enterprise operating system and applications such as SQL server or Exchange Server on both the live server and the recovery server. If you were previously running standard editions there is a significant amount of work and disruption required in the live environment before we can even think of introducing the protection we require.

The levels of cost and administration time involved typically meant that companies selected only the most business critical applications to be protected, accepting less or no protection other than system backups for the remaining systems.

Even after all of the above work has been completed no one wanted to hear the following word “testing”. Testing historically meant out of hours working and the real possibility of data corruption; failover to the recovery node was often straight forwards, however failback from the recovery node to the live server was often not.

In my experience the vast majority of people could not confidently say their IT systems disaster recovery process could recover the systems within the Recovery Time and Recovery Point objectives stipulated by the business. In some cases some people have admitted they simply don't know if it works as the process had never been tested due to the amount of work and risk involved.

## Enter the virtual world

So what are the key differences in a virtual server environment?

### *Application & data protection*

Due to the nature of a virtual server environment we protect the entire virtual server, not just an individual application running on the server. We clone and replicate the entire live server including all of the installed applications and their associated data. These replicated virtual servers are kept up to date via a bespoke replication schedule.

We have a number of DR solutions that are capable of supporting both physical and virtual production servers and we can tailor the replication schedule and solution to meet our clients' specific requirements.

The hardware used at a recovery site can be vastly different from the hardware on which the live systems are running, different vendor; different model, lower specification hardware can be used on which to recover the failed live server(s). Additionally the consolidation ratio of virtual machines on a single physical platform means that the 1:1 hardware correlation which exists in the physical server environment is removed.

What does this all mean? Well, a single server such as a DELL PE2950 that has been replaced by newer hardware could have its memory, CPU and disk resources increased and could be used to protect and recover 5-10 live virtual business servers.

### *Testing*

The painful part, or it was! We demonstrate to clients on a regular basis the recovery of our internal systems at SITS Group. We show the entire recovery process from start to finish, depending on the size of the virtual server being recovered this takes around 10 minutes.

Once the server has been recovered we login, we view the data on the recovered server, and we validate the time stamps on the recovered files and confirm the integrity of the data. We do all of this whilst the live server remains up and running and in use by our staff.

Once we have finished the recovery test we simply resume the protection of the live server and the data replication schedule continues ensuring the recovery data is up to date.

We know our disaster recovery processes work, we test them a couple times a week whilst our clients or potential clients watch (nothing like pressure). When did you last test your disaster recovery processes?

### *Failback*

So having recovered a failed live system to a recovery virtual server and run on the recovery server for a period of time, how do we failback to the live systems? The answer to this question traditionally has been "not easily"; in a traditional

world, the process of failback from the DR site to the production site involves the same risks, manual effort and downtime as the original failure.

Having a virtual infrastructure however, makes this process both painless and streamlined. Our replication solutions provide the facility for one button failback ensuring that the entire DR process is easy to manage and effective.

### *Summary*

In short, virtualisation technology enables more businesses to implement cost effective disaster recovery strategies, and arguably more importantly, it provides confidence to those companies that their DR strategies work. As well as allowing for easy testing of processes with greatly reduced numbers of people, time and risk involved.

For high level details of a recent DR solution deployed by SITS Group click on the following link:

<http://bdaily.info/news/technology-and-science/19-11-2009/technology-proves-virtually-no-risk/>

See how our DR solutions can streamline your DR processes and provide your organisation with a robust solution which can be easily tested, audited and managed by contacting sales on 0191 2155015 or email [enquiries@sitsgroup.com](mailto:enquiries@sitsgroup.com)