

HOME | CURRENT ISSUE

Best Practices: Foolproof DR is still a moving target

http://searchStorage.techtarget.com/magazineFeature/0,296894,sid5_gci1305544,00.html

by: James Damoulakis

Issue: Mar 2008

Keep these eight key steps in mind when designing and testing your disaster recovery strategy.

In a recent conversation I had regarding disaster recovery (DR), a CIO remarked that he'd like to achieve what he called "provable" DR. I heard this as evidence of a positive development I've noticed in recent years: Many companies have finally become serious about DR. There are several reasons for this, including the heightened awareness of data protection issues among the general public.

Technology developments have also played a critical role, with more data protection options available than ever before. The changing nature of business applications means that expectations regarding performance and availability have also been raised.

But achieving DR "provability"--or at least greater predictability--remains a challenge. Fundamentally, DR is a holistic endeavor with a number of moving parts. It's fairly easy to deal with one component of DR and for it to perform reasonably well. The hard part is ensuring the coordination and synchronization of the various elements so they function together. To establish more predictable DR, I've outlined the following eight necessary elements.

1. Clearly defined organizational responsibilities. Roles and responsibilities is a major area where organizations fall short with regard to DR. The DR process is much more than restoring or replicating data; it's about ensuring that applications and the systems they support can be returned to functional business usage. Accomplishing this requires participation from groups outside of IT, including corporate governance and oversight groups, finance and the business units impacted.

While IT may drive the planning and execution of DR, it's imperative that it's coordinated with a broader business-continuity planning effort. A solid DR strategy needs the strong endorsement of the highest level executives.

2. Validate the business impact analysis (BIA) process. Technically, the BIA isn't part of the DR process--it's a prerequisite that forms the foundation of DR planning. In a perfect world, the output of a BIA would define the kinds of recovery capabilities IT must design and deliver in support of the business. The real world, unfortunately, isn't so simple. Information is often incomplete, and we need to make assumptions to fill in the gaps.

However, there needs to be a level of confidence in the validity of BIA requirements. To design an effective DR strategy, IT needs two pieces of data: recovery requirements (e.g., RTO and RPO), and a reliable estimate of the real cost of downtime to the business. Recovery requirements supplied without valid financial impact data can result in a great deal of effort for a project that isn't funded. Passing requirements through the financial prism adds realism to the process.

3. Define and tier application recovery services. When business executives hear IT people talking DR strategy, they're thinking cost. DR represents insurance and because no one wants to buy too much insurance, efficiency is vital. While there are significant fixed costs inherent to DR--a recovery site, for example--there are also a substantial number of variable costs that can be controlled. The key is to realize that not every application requires a two-hour recovery time. Establishing a catalog of services based on BIA requirements that provide several levels of recovery and then aligning applications appropriately is one way to contain costs. With multilevel

recovery services, applications can be prioritized according to importance. Among the business attributes that should be defined within the service catalog are risk (usually expressed in terms of RTO and RPO), quality of service (including performance and consistency levels) and cost.

4. Implement a comprehensive cost model. While the BIA determines the impact of downtime to a line of business, and tiered recovery services provide a catalog of services that align with business requirements, there also needs to be a method to determine and allocate the cost of those services. Corporate governance may help set thresholds for recovery and imply minimum levels of protection, but the service level is greatly influenced by cost. The cost model should calculate the per-unit total cost of ownership that would be charged to the business for any given service offering. Among the items included in such a cost model are personnel, facilities, hardware and software, maintenance and support. Having this data available helps significantly in aligning "want" with "need," and is a critical success factor in delivering these services efficiently.

5. Design an effective DR infrastructure. The DR infrastructure must support the BIA requirements and service-level targets. While DR is an extension of operational recovery capability, factors such as distance and bandwidth also come into play. The good news is that the number of remote recovery options available to architects and designers has increased dramatically over the past few years. Traditional storage mirroring and replication are more broadly available on a wide range of systems, and compression and deduplication technologies can reduce bandwidth requirements. In addition, technologies like server virtualization can dramatically improve remote recoverability.

6. Select the right target recovery site. DR site selection often presents a challenge. Organizations with multiple data centers can develop cross-site recovery capabilities; if you don't have that option, selecting a DR site can easily become the biggest challenge in getting DR off the ground.

Key concerns include the levels of protection needed, and whether to own or outsource (and to what degree). The two chief (and often competing) factors to consider are risk and convenience. Planning for protection against a regional disaster means that many DR sites get pushed far away from headquarters, where most of the IT staff is housed. Service recovery levels will determine whether the site is a "hot," "warm" or "cold" site. This is a critical designation because there's a substantial difference in the fixed cost of each. Generally, RTOs of less than a day require a hot site. The question of outsourcing depends on the desired degree of control, guarantees of infrastructure availability at a given location and, of course, cost.

7. Establish mature operational disciplines. One of my colleagues is fond of pointing out that one of the best ways to improve DR is to improve production. Put another way, if normal day-to-day operations don't tend to function well, your DR isn't likely to either. Therefore, operational discipline is an essential element of predictable DR. The first sign of a potential operational deficiency is the lack of documentation for key processes. Given that DR, by definition, occurs under seriously sub-optimal conditions, the need for well-documented standard operating procedures is clear. Organizations that have established and actively embrace standard frameworks, like the Information Technology Infrastructure Library (ITIL), are significantly improving their odds of recoverability in the chaotic atmosphere of a disaster situation.

8. Develop a realistic testing methodology. Given the operational disruption, practical difficulties and costs involved, we tend to focus our testing on those components that are easy to test. But realistic testing is just that--testing real business function recovery. While it's necessary to perform component testing on a regular basis, it's equally important to test the recoverability of large-scale functions to ensure that interoperability and interdependency issues are addressed. The closer to a real production environment a test can get, the more "provable" the DR capability.

The elements outlined here transcend the boundaries of the IT infrastructure. It's therefore critical for IT administrators to have a strong understanding of the problems at hand and to learn how to

address them so they can influence strategic decision-making wherever possible. This will help them avoid being placed in the Catch-22 situation of solving a problem over which they have no control.