Will your current DR plans truly hold up in the event of a disaster?

CONSIDERATIONS CHECKLIST

SUNGARD®
AVAILABILITY
SERVICES™

Executive Summary

Cost-cutting pressures on CIOs have led many to accept more risk and settle into a mindset that their Disaster Recovery (DR) program is comprehensive, when, in fact, it might not hold up when matched against the company's business requirements.

Quite often, DR planning is not top of mind for senior management. With shrinking IT budgets and teams, keeping up with day-to-day responsibilities is challenging, and many IT teams will tend to focus on only the infrastructure, hardware, and software, while neglecting the people and processes that are needed to execute the plan. Introducing complexity into an IT business practice that is under-resourced can lead to an incomplete DR plan, failed recovery testing, or even worse, a company that has declared a disaster without any chance of recovering quickly.

To avoid failure, it is important to truly understand how your company addresses recovery management. In some instances, your IT teams may be reluctant to publicly identify areas of concern, while some may not even know whether their DR plan will work.

This checklist reveals 10 truths about DR that IT teams are most likely reluctant to tell their CIOs, and provides you with questions you can use to ask your team about your recovery management strategy.

1

We are not able to meet our RTO/RPOs for our mission-critical applications.

Maybe you passed your last annual DR test, or maybe you didn't. Even if you did pass, the test is only a predictor of whether or not you are actually able to meet the Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) dictated by your business requirements. What many IT leaders do not consider is that DR tests are conducted under managed conditions, and can take months to plan.

Most causes of outages (power failure, human error, hardware failure) do not give you notice. The single most important factor in determining whether or not your recovery management plan is successful is in your ability to reflect the day-to-day change management tasks so that they are perfectly in sync with your production environment. Today's mission-critical applications have many dependencies that change frequently. Without ongoing tests, a recovery plan that worked before might now fail to restore availability to a vital business application.

Question:

When was the last time we successfully tested all mission-critical applications against our RPO and RTO measurements?

2

Our DR plan just scratches the surface.

You need to consider your recovery management capabilities in the context of the impact it has on staff and the long-term availability of your datacenter. Determining how long you can support an outage at your recovery center should impact your DR plan approach. It is also important to understand how the secondary site will be managed. Mostly likely, you will need to send staff to the secondary site to work on the recovery and maintain the temporary production environment, but this may not be easy in the event of a natural disaster. You cannot assume that the right people can get to the right places for each identified disaster.

Provided you have the capabilities to recover, you must ensure that your organization is well informed of procedures and chains of command. Such issues came to the forefront after Hurricane Sandy, when massive flooding kept many roads closed for weeks, and gasoline shortages prevented many people from getting around at all.

Question:

What would we do in a major disaster if we lost power for days or weeks, lost buildings, or lost communications links?

3

We know how to failover to a recovery site, but we lack the experience and capabilities to know how to failback.

Failover and failback are critical to executing a DR plan. Failback can often be the most disruptive element to DR execution. With failback, most processes must be reversed. When a failover occurs, the secondary backup site must be a duplicate of your primary site. It must be able to support your production environment and offer the same protections needed to function as the primary site for a period of time. Failback means that your organization is looking to reinstate the production environment. Recovering back to the primary environment works the same way as a failover except in the opposite direction. Testing this scenario should also be performed, documented, and controlled. Not documenting and testing this component of your DR plan could force you to rely on your secondary site for extended periods of time, adding significant cost to the business (not to mention a probable strain to your staff).

Question:

Do we test our capabilities to failback during our schedule recovery tests?

4

Our runbooks are probably unusable.

Your runbooks should contain all of the information you and your staff need to perform day-to-day operations and to respond to emergency situations, including resource information about the primary datacenter and its hardware and software. Step-by-step recovery procedures for operational processes are also a critical component. If the procedures are not frequently updated, or not thoroughly vetted with key stakeholders, your recovery process will be significantly slowed, if not outright halted. And remember, the more time it takes to recover, the more expensive it gets. The Aberdeen Group estimates that downtime cost the average company $160,000 per hour in 2012.[1]

Question:

How often do we evaluate and update our DR plan?

[1] "Working in the dark: the financial impact of IT downtime," TechJournal, December 7, 2011.

5

We haven't changed when it comes to change management.

With today's highly dynamic production environments, change is constant. Next generation datacenter technologies such as virtualization make it easier to create and deploy applications, allocate and provision storage, and set up new systems. However, the ease and frequency at which these changes occur can prevent your team from properly recording them at your recovery site. Without properly performing change management, secondary and backup environments can quickly get out of step with your production environment, causing recovery failures.

Question:

What are we doing to ensure that our testing environment reflects our live production environment?

6

We can pass an audit, but it doesn’t mean we’re recoverable

Passing an audit means you have a plan that meets a specific list of requirements. It does not mean that your plan will provide recoverability. Most auditors do not focus on the variables of your DR plan, and do not look at the effectiveness of the plan for each and every disaster scenario. They only ensure that you have met the static requirements established in the audit itself. The fact is, you can pass an audit, but still fail to recover from an actual event.

Question:

When was the last time we tested the DR plans?

7

Our IT environment is getting too complex.

Your business environment is becoming more dynamic, and becoming dependent upon an increasing number of applications. A critical Tier-1 application might need a database that you have classified as Tier-3 in order to run properly. Restoration of full services will require recovery of all of these elements. As a result, you will need to tier your applications accordingly, which may require adjustments to your tiered environment to ensure you are addressing all interdependencies. A complex infrastructure will make the tiering — and therefore, the recovery — that much more difficult.

Question:

How have we tiered our applications to aid in recovery?

8

Backing up doesn’t move us forward.

Backing up is not, by itself, a DR solution, however it is a critical component to a successful recovery management plan. Whether you are replicating data to disk, tape, or a combination of both, moving data between storage mediums is s-l-o-w. If it takes an unacceptable amount of time to move and restore data, then testing is probably out of the question. Time-to-restore concerns have also led companies to forego a regular test restoration process, which can lead to lost data.

Question:

How have we integrated data management into our recovery management program and testing?

9

We're neither testing enough nor do we have the time or people to do it right.

Only 20–30 percent of BCDR plans are tested, and many of those fail.[2] While you have a DR plan and can arguably restore most of your mission-critical applications through testing, measuring the total restore against the frequency of testing and mapping out the resources needed to conduct and validate a test must also be performed. You might have the plan in place, but without the resources available or conducting the actual test, you cannot validate success. Testing recovery procedures of applications is a lot different than recreating a datacenter from scratch, and a 72-hour testing window is not adequate; it is just enough time to corral the right employees and beg them to participate in the test when it is not a part of their core function. Companies will often work with whatever resources they have. In professional terms, they "wing it."

Question:

Do we have the bandwidth and expertise in-house to be fully recoverable?

[2] Sungard Availability Services, 2012. How Sungard AS is Changing BCDR Testing, an Enterprise Strategy Group Product Brief.

10

We’re not keen on the idea of someone else doing our recovery.

Your employees are likely aware of the outsourcing option, which could lead to fear about their own job security and, on a more practical level, about how it impacts their level of control.[3] However, the benefits of partnering with a recovery service provider actually complements their skill sets by allowing them to focus on strategic projects rather than operational tasks, while improving the overall recoverability for the business. According to a recent IDC study, enterprises that performed IT recovery in-house lost an average $4 million per disaster incident across a variety of business functions (e.g., sales/marketing, financing, e-commerce), while enterprises that partnered with a managed recovery service provider lost an average of $1.1 million per incident.[4]

Question:

Do we have the recovery expertise in place to ensure that your role can be successful?

[3] Outsourcing Disaster Recovery Services vs. In-House Disaster Recovery.

[4] IDC Study

Conclusion

Simply having a DR plan in place does not guarantee recovery from an outage or disruption. Backup plans need to be complemented with well thought-out processes, and staff must be well-trained in carrying out the needed procedures.

Additionally, the dynamic nature of business compute environments means that DR plans need to be frequently evaluated, updated, and tested. Change management must be taken into account to ensure the secondary systems are aligned with those at the primary site.

Because of the critical need for availability in today's competitive marketplace, you do not have the luxury of trusting that proposed DR strategies will work as advertised. You must ask the right questions to get at the heart of the matter. Will your current plans truly hold up in the event of a disaster?

For more information visit the **Managed Recovery Program** web page.

Sungard AS Managed Recovery Program
Sungard AS Managed Recovery Program (MRP) is an integrated approach to IT disaster recovery management that provides IT organizations with the ability to focus on production, improves overall recoverability, ensures a constant state of readiness, and reduces costs. By selectively outsourcing the "DR program" elements to Sungard AS — recovery procedure development, test planning and execution, execution of recovery procedures at time of test and disaster, and ongoing life cycle and change management — IT can improve DR efficiency and effectiveness, reduce complexity, and repurpose precious staff availability away from non-core processes and towards business value creation.

Additional reading

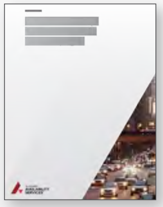

Managed Recovery Program Brochure

Disaster Recovery Service Provider Checklist

About Sungard Availability Services
Sungard Availability Services provides disaster recovery services, managed IT services, information availability consulting services and business continuity management software.

To learn more, visit **www.sungardas.com**
or call 1-888-270-3657

Trademark information
Sungard Availability Services is a trademark of SunGard Data Systems Inc. or its affiliate used under license. All other trade names are trademarks or registered trademarks of their respective holders.

Connect with Us