Disaster Recovery in the Cloud

Quick Guide

This quick, action-oriented guide will help you set a cloud-based disaster recovery (DR), also known as Recovery-as-a-Service (RaaS), plan in motion. The document will assist IT directors and managers in creating and implementing the most effective DR strategy for their organization. If you are looking for additional background information on Recovery-as-a-Service, learn more with The Business Case for a Multi-Tenant, Cloud-Based Recovery-as-a-Service Solution.

Checklist

❏ **List your applications**
A quick export and sort of your VM list from VMware vSphere® is a simple start.

❏ **Perform a risk assessment and business impact analysis for each application**
Rank each application 1–5 on how critical it is to daily operations.

❏ **Determine each application's ideal RTO and RPO**
Relying on the risk assessment, think about how long the business can function without the application and how much data (time) can be lost in the worst case.

❏ **Determine your Disaster Recovery and Backup budgets**
Your budget may not allow for each application to be protected with minimal RTO and RPOs. Establish a budget framework for your solutions.

❏ **Prioritize your applications**
Using the data from steps two and three, group your applications into recovery tiers that you can spread your budget over.

❏ **Determine your recovery destination**
Review Recovery-as-a-Service providers to find one that meets your budget and recovery demands.

❏ **Set your runbook**
Whiteboard the recovery process step-by-step to identify and document the boot order of your VMs and any other changes that may need to be made while recovering. Consolidate the steps into a Recovery Procedure runbook.

❏ **Plan and practice the failback process**
Actually test it!

❏ **Plan for the unexpected**
Keep runbook and recovery contact information where you can get to it, even if you don't have access to any of your other systems like email or LAN.

❏ **Establish seasonal testing schedule**
Test, test and then test again.

Step Details

1. List your applications

You can't protect your applications if you don't know where each application is and what its core responsibilities are. Particularly if you're a member of a large enterprise, it's likely each individual business unit is better in tune with its applications' needs and usage and what will be most important for each unit to keep functioning in the event of a declaration.

A quick export and sort of your VM list from VMware vSphere is a simple start. An ongoing approach to maintain a full inventory is to start a running dialogue with each business unit from the very beginning of this process. You'll be conducting more in-depth interviews during the next stage, so starting out the conversations by inventorying all applications is a great, neutral place to start the conversation.

If possible, establish one point-person within each business unit who will be responsible for maintaining contact with you. As things change and develop over time, and they will, you will have a running dialogue with this person to discuss new applications that are being phased out and any significant changes. This ensures you aren't protecting legacy applications or leaving new applications unprotected.

2. Perform a risk assessment and business impact analysis for each application

Once inventoried, you will need to work in conjunction with the business unit to perform a risk assessment and business impact analysis for each application in your company. This will be a time-consuming process determined by the size of your organization, but it's vitally important to making sure your plan protects what is most important to the business.

When conducting your risk assessment and business impact analysis, remember to look at the totality and connectivity of the applications. The business unit may not be aware of how and to what the application is connected to, so it's your job to keep an eye on the connectivity factor.

You should also consider not just outage scenarios, but disruption and interruptions that could occur intermittently and impact the business as well. It might not be the hurricane that hits your datacenter, but rather the failed or faulty wire that cuts in and out of service.

When evaluating the business critical nature of an application, look beyond just the revenue-generating applications and remember to include customer service, financial and other core functioning applications that will be necessary in the event of a disruption.

3. Determine each application's ideal RTO and RPO

The main function of your business impact analysis and risk assessment is to determine the ideal RTO and RPO needs of each application.

It’s not just about the critical importance of an application to the business overall, but your assessment should be about the time-criticality of each application. Consider how many and how often the application processes transactions and new data, as well as how soon the application would need to be up and running after a disruption.

RTO stands for Recovery Time Objective. The Recovery Time Objective is the ideal time it would take for the application to be stood back up after a declaration. This may be seconds, minutes, hours or days depending on its time sensitivity and impact on the business’s ability to recover from a disruption.

RPO stands for Recovery Point Objective. The RPO will be determined by a combination of the application’s rate of change and the nature of the application’s business purpose. RPO refers to the point in time in the past from which data will be recovered. For example, if the application has an RPO of 30 minutes, the 30 minutes of data between the RPO and the disaster event will not be recovered, but everything prior to that point will be.

4. Determine your disaster recovery and backup budgets

If your organization is like most others, the applications will each come back with assessments requesting minimum RTO and RPOs. That’s why the next step is to determine what budget you have to spread over the applications.

Once you’ve determined your total budget for DR, recovery and backups it is time to prioritize your applications in order of time-sensitivity to recovery. There may only be a few applications total that will warrant the budget to recover with the minimum RTO and RPOs. Some applications may be fine with a downtime of 3-5 days on backups and some applications may be considered test/dev and not warrant protection at all.

A combination of the budget and a prioritization list will help you sort which applications need to be addressed first and which can wait a little longer.

5. Determine your recovery destination

Perhaps your company is large enough to have private datacenters in multiple locations. If that is the case, consider replicating to your secondary sites.

The site to which you will recover is also greatly dependent on the type of technology you decide to use to protect and recover your application. Some applications may require a “hot-site” approach, sometimes called an active/active approach in which there is no downtime. This is the most expensive protection method, so it should be used for only the most critical applications.

Others may require a “warm-site” approach, which could include cloud-based RaaS solutions. With RaaS, applications are replicated to a second site near continuously. RaaS can be one of the most economically efficient solutions because it affords an ideal mix of RPO and RTO promises and because you’re only paying for storage and bandwidth before you failover, the cost is a fraction of traditional solutions.

Cloud-based RaaS is an exciting addition to the DR space because it allows clients the peace of mind that wasn’t available before. Cloud-based DR users can see the replication working, click and test their solution at any time and they have the backing of a service provider in case things get tricky. The price point is compelling because it can often beat legacy prices with hard equipment. For a warm-site approach, it's often the best and most compelling option.

Some applications can withstand a longer downtime and would typically require a "cold-site" approach. This includes backup-to-tape and cloud backup products. This option protects the data, but not the applications themselves.

6. Plan and practice the failback process

It's easy to focus on the failover and forget one of the most important parts of the entire process: the failback.

Failback can be harder than failover. In failover you're often failing over to a clean environment that was previously unused. In failback you're contending with an environment that has been running and could have been impacted by your disruption. Many DR products simply don't have a failback option and the process may be more disruptive than the initial incident.

Be sure to determine the appropriate failback strategy and write your plan accordingly. Failback should be tested too, so you don't find yourself having recovered from one disruption only to have a second, bigger problem when you realize you can't failback properly.

7. Test, test and then test again

The value of your solution is only equal to your confidence in your ability to recover. With legacy recovery technology, tests were nearly impossible to perform so you were stuck crossing your fingers and hoping it worked.

A common misconception is that you only need to test once successfully and then you're done. Your environment is a living thing that changes over time in ways you can't always predict. You should test at least twice a year, if not more, to ensure your solution is always up-to-date. A static DR plan will likely be less successful in recovery compared to one that is tested at least bi-annually. Testing is not just about making sure a proper failover will happen, but it's also about identifying and compensating for changes that have happened since your last test.

If you have a test that finds errors, that is a successful test because it means you've established errors that you might have run into during a real disaster.

When comparing solutions, testing with cloud-based RaaS is affordable and easy to execute compared to legacy technology. It can often be pre-scheduled and done with the click of a few buttons. Some vendors even reimburse the resources used during pre-scheduled tests because they know the benefits of testing outweigh anything else you can do to prepare for a disaster.

When you're setting up your DR plan, also consider planning for the unexpected. Your plan may rely on your staff to execute the plan, but in the event of a disaster, they may not all be available. If you work with an external provider to help you plan and execute, you can avoid the worry of not having staff available and rely on your provider.

You should also think specifically about applications that may be enterprise-wide apps and not owned by specific departments, or that are owned by the datacenter. These can easily fall through the cracks and can be the most important to protect.

Now that you have your checklist in place, get started! If you would like help assessing your risk or establishing your DR plan, particularly with cloud, Bluelock is available and ready to help. The Professional and Advisory Services, as well as the Cloud Support team, have seen it all and can help you create a plan to protect and recover your applications regardless of what strikes.

For more information on Bluelock's leading RaaS solutions or to speak to an expert visit **www.bluelock.com/cloud-services/raas** or call **888-402-2583**.