



White Paper

Step-by-Step Planning for Disaster Recovery

- Introduction 3
- Phase One: Strategy 3
 - Step 1: Define Recovery Objectives** 3
 - Step 2: Know Your Response Team** 3
 - Step 3: Set Up Lines of Communication** 4
 - Step 4: Document Your Resources** 4
 - Step 5: Select a DR Location and Minimize the Costs** 4
- Phase Two: How to Take Action 5
 - Step 1: Know When to React** 5
 - Step 2: Know How to React** 5
 - Step 3: Getting Back to Normal Operation** 5
- Phase Three: Testing 6
 - Step 1: Test Everything, Then Repeat** 6
 - Step 2: Update Routines Regularly** 6
 - Step 3: Stay Vigilant** 6
- Disaster Recovery Template Plan 7
 - Disaster Recovery Plan Stages** 7
 - System Definitions 7
 - Incident Detection and Response 7
 - System Restoration 7
- Summary 8

Introduction

Disaster recovery remains the most pressing concern for businesses in the cloud today. Although ransomware is a chief concern, additional threat vectors such as human error, natural disaster, network failures, and software vulnerabilities make developing a disaster recovery strategy the top IT priority for every business today, whether it operates in the cloud or not.

This white paper discusses a three-phase process companies can use to begin to conceive, design, and enact secure disaster recovery plans. Starting from the initial strategy sessions, through the design phase, to testing the solutions you put in place, this guide gives step-by-step instruction to making sure of business continuity when the worst takes place.

Phase One: Strategy

Your disaster recovery plan starts with building a thorough strategy. Here you discover the limits of your company's ability to tolerate a disaster event, lay out and document all of the goals and procedures for your company's response to that disaster, and establish the tools you use to get your operation back online.

1 Step 1: Define Recovery Objectives

Your overall operation has a bottom line for how long it takes to recover from an attack or failure before it begins to cause harm. There are two measurements for disaster recovery that are universal to business continuity: these are your deployment's recovery point objective (RPO) and recovery time objective (RTO).

Your RPO is the length of time that your business can withstand a loss of data before full recovery. Inability to access after the RPO is considered to be harmful to the business and its ability to operate. The RTO is the length of time it takes from when a system is affected to full system recovery. A company must be aware how long it can tolerate a loss of function without severe loss of business, traffic, or SLA. In some cases there can be not only financial losses but also legal repercussions should RTO be missed.

2 Step 2: Know Your Response Team

In this second step of the strategy phase, you have to identify all of the individuals who make up your disaster recovery response personnel. This list should include obvious team leaders such as the CEO and CTO (and possibly the CFO and COO, depending on your company). It should also extend to the engineers and system administrators who are instrumental to bringing the system back to normal.

Additional team members to include in the response team are the public relations department head and the contact persons at any major customers, partners, or vendors who might rely on your company for their normal business operation. This list of personnel should also—to the extent that anything that happens in a disaster is predictable—outline some of the expected tasks those individuals would be expected to perform in a crisis.

3 Step 3: Set Up Lines of Communication

After personnel have been identified and their roles in course of a disaster and its recovery are established, the next step is to establish clear lines of communication, which are followed during the disaster.

You should also be prepared to notify the proper authorities to contact in case of a malicious attack or other failure.

Specific communication applications should be selected to use in the disaster scenario. In addition to the electronic channels of communication, a physical site should be designated for face-to-face crisis response conferencing. Another crisis communication hack is to store backup cell phones—with sufficient plan coverage and battery charge—that can be used for phone calls or SMS, because normal communications over your network might be disrupted.

4 Step 4: Document Your Resources

The next part to drawing up your plan is to document your resources. Resources to list in this document include all of the network connections, all of your system's requirements, your system's rate of usage, OS system details, running applications, and the location of all stored data and backup data. In short: you want every aspect of your IT deployment to be laid out in full. This information is no small amount, but it is crucial in a disaster event.

5 Step 5: Select a DR Location and Minimize the Costs

Whether you deploy in the cloud or still have your data center in the building with you, choosing the right disaster recovery tools is an important part of the plan-making process. The cloud, however, has a few resources available that make it the more attractive choice for disaster recovery.

Using the resources provided by cloud providers such as AWS or Azure, backup sites can easily be run using cloud provider resources such as the [AWS pilot light architecture](#), giving you a failover site that can be ready to act instantaneously. The big cloud providers also offer automated disaster recovery as a service (DRaaS), including [Azure Site Recovery](#), ready to act at any time with orchestrated recovery response to keep your company safe.

With so many elements in play for a disaster recovery, your total cost of operation might become a point of concern at this stage. Costs are a significant challenge when maintaining effective disaster recovery solutions.

One way to manage costs is to select tools that reduce the amount of storage space or bandwidth it takes to maintain backup data. NetApp offers solutions to both those problems: NetApp® [ONTAP® Cloud](#) can help reinforce your company's disaster recovery capabilities by offering time-saving and cost-cutting storage efficiencies, and [Cloud Sync](#) makes [transferring files to and from the cloud and any NFS/CIFS share](#) seamless and fast.

ONTAP Cloud acts to save on storage space by using powerful features such as [tiering infrequently used data to Amazon S3](#), data compression and deduplication, and thin provisioning. Combined, these features utilize storage more effectively, which creates a significant savings for keeping up-to-date copies of your business-critical data. Cloud Sync uses parallel processing to make it possible to keep data synced [more efficiently than any homegrown migration tool](#) you might create on your own with a resource such as rsync.

Phase Two: How to Take Action

In this phase of your disaster recovery planning, you design the actions that you take in the course of a disaster event. A useful document to create based on the steps outlined in this phase is a flowchart outlining the cascading reactions to the initial threat detection. This flowchart should be appended to the disaster recovery information that you compiled in the strategy phase.

1 Step 1: Know When to React

This step is where everything hinges. How can you tell if an incident can be easily addressed or if it is going to cause harm to the business?

One best practice for knowing when to react is to employ an alert and monitoring system that can notify you as soon as red flags are raised. Another useful resource for determining disaster activity is from your business's users themselves: when they can't access the site, you'll know about it.

When it seems as if an event might be taking place, the first steps to take should be to assess all of your server hardware, cloud service, logs, network, and anything else IT-related to look for problems. If any one of these components isn't operating normally, it is an indication that it is time to put your disaster recovery plan into effect.

2 Step 2: Know How to React

When it's clear that what you're dealing with is in fact a disaster, immediately activate your backup failover site. Booting up the backup site is not the end of this step. There is a lot of verification that has to be done to make sure that it can operate as close to the normal production environment as possible. That means verifying that data is synchronized and the operation of critical site features. These processes and the teams responsible for them should all be listed and labeled in the disaster recovery plan document.

3 Step 3: Getting Back to Normal Operation

Backup sites aren't meant for continuous deployment because they aren't designed for a full production workload; the sooner you can get back to the primary site, the better. This process could result in loss of services as key parts of the deployment failback.

Anticipate the order in which elements of your application will return to normal operation and have an estimated time that the failback will take, including any anticipated loss of services. Some cloud providers have solutions that offer automatic failback when the primary site is back online; this solution can be helpful in making sure that there is no break in services.

Phase Three: Testing

Things might have bounced back to normal, but that doesn't mean you can rest easy. After you know how your disaster recovery plan will react to a catastrophic event, you need to design the part of your plan that tries to make sure that it doesn't happen again—or ever. A big part of that means testing, but it also means staying on top of updates and paying constant attention to your deployment.

1 Step 1: Test Everything, Then Repeat

Every part of your disaster recovery plan should be fault tested. After that happens, test it again. If you don't test your plan, it might not work.

Failover tests to backup sites should be planned to take place **at least once or twice a year**.

Not every company can afford it, but **the ideal is to perform failover tests once every quarter**. These tests won't just give you an assessment of the health of your sites. They also give you a chance to see your personnel in action and make any necessary changes to the structure of the disaster recovery plan that was outlined in your strategy phase. Because tests are designed and run to meet individual deployment needs, they are mostly designed internally. One testing resource for deployments on AWS is [Netflix Simian Army](#).

These tests don't always need to be full-on wargame scenarios. Testing can be integrated with the normal operation of your business. Shifting less-important workloads to the backup site in order to offload usage on the primary site is not out of the ordinary; such shifting gives administrators a good idea of how well the backup site holds up to usage.

2 Step 2: Update Routines Regularly

Testing gives an idea of how your disaster recovery system will perform in a crisis, but you need those elements to run at the latest build possible to make sure you aren't vulnerable to the latest threats. Patching and updating should all be scheduled to automatically bring your system up to date with the latest fixes.

Another important area to keep up to date is your disaster recovery plan documentation. These documents should have details about how every aspect of your disaster recovery plan behaved during each test and give accurate assessments for the performance of team members. This information should be used to adequately adapt to any weaknesses that might have been discovered during the test or actual disaster event.

3 Step 3: Stay Vigilant

Even after all the previous steps have been taken, there is always a reason to worry about the vulnerability of your deployment. This situation is not necessarily a bad thing; one of the biggest mistakes a company can make when it comes to cybersecurity and its disaster recovery preparedness is negligence. Remaining vigilant against all possible threats is a sure way to remain one step ahead of the next attack.

Ways to stay vigilant include employing threat-monitoring systems that can keep an eye on your primary and failover sites. Using a cloud service provider makes it possible to take advantage of additional tools for monitoring, such as [Azure Security Center](#) and [AWS CloudWatch](#).

Disaster Recovery Template Plan

Although the following template can't fully simulate the actual length of the average company's disaster recovery plan document—something that can span as many as 50 pages—this template gives you an example of the kind of information that your plan should include in a layout you can use.

Disaster Recovery Plan Stages

	System Definitions	Incident Detection and Response	System Restoration
Time frame	Months to years before incident occurs	Up to 24 hours from initial detection until response takes effect	Up to 36 hours from initial detection until system is fully restored
Action items	<p>Determine your RPO and RTO points.</p> <p>Draw up definitive steps for disaster recovery procedures.</p>	<p>Notification of threat by monitoring service or complaint.</p> <p>↓</p> <p>Response team deployment.</p> <p>↓</p> <p>Analyze situation to make the call if restoration can be achieved before RTO or if disaster recovery steps should be taken. If disaster recovery is put in motion.</p> <p>↓</p> <p>Activate the backup site. Ops should take all necessary steps to restore VMs, telecommunications, SAN.</p> <p>↓</p> <p>Validate data integrity of database, looking for potential data loss.</p> <p>↓</p> <p>Reroute all customers to failover site.</p>	<p>After the disaster has been addressed, move to restoration.</p> <p>↓</p> <p>Verify that normal site operations are functional.</p> <p>↓</p> <p>Check network, application, and other component services.</p> <p>↓</p> <p>Test use of primary site before returning services to there.</p>

	System Definitions	Incident Detection and Response	System Restoration
Reference material	Disaster recovery strategy document (sections on RTO/RPO information, response personnel contact lists, infrastructure register)	Disaster recovery strategy document (sections on infrastructure details, failover steps, cloud service support details, if any) Response and resolution flowchart	Disaster recovery strategy document (infrastructure details, failback steps) Response and resolution flowchart

Summary

The three phases to building your disaster recovery plan given in this document should offer the essential outline you need to design an effective way to recover from debilitating events. Although this guide is a good way to start, it isn't a one-size-fits-all solution. You have to tailor your company's specific response to your specific needs.

No matter how much you value the full operation of your business, a skyrocketing budget can cause hesitation with fully investing in an appropriate disaster recovery plan. Turning to the cloud-based disaster recovery solutions offered by [AWS](#), [Google Cloud](#), and [Azure](#) is one way to reduce the costs of on-premises disaster infrastructure, but even that might not be enough.

The best approach in that scenario is to turn to trusted third-party vendors for further cost-cutting and disaster recovery capabilities in the cloud, such as NetApp ONTAP Cloud and Cloud Sync. Efficient RTO and RPO require full backups to be stored at all times. ONTAP Cloud provides you with a way to keep storage costs for that data under control through the use of [powerful space-saving features](#), and Cloud Sync makes migration easy and efficient.

If you are ready to take the next step in creating a disaster recovery plan that is both effective and cost-conscious, NetApp can help you with a free 30-day trial of ONTAP Cloud [on Azure](#) and [on AWS](#) and with a [free 14-day trial of Cloud Sync](#).

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners

Copyright Notice

External Publication of NetApp—Any information that is to be used in advertising, press releases, or promotional materials requires prior written approval from NetApp's CEO. A draft of the proposed document should accompany any such request. NetApp reserves the right to deny approval of external usage for any reason. Copyright 2017 NetApp. Reproduction without written permission is completely forbidden.