Business Disaster Recovery Strategies: Succeeding Where the Odds Say You Will Fail

Michael White
Global Technical Evangelist, Veeam

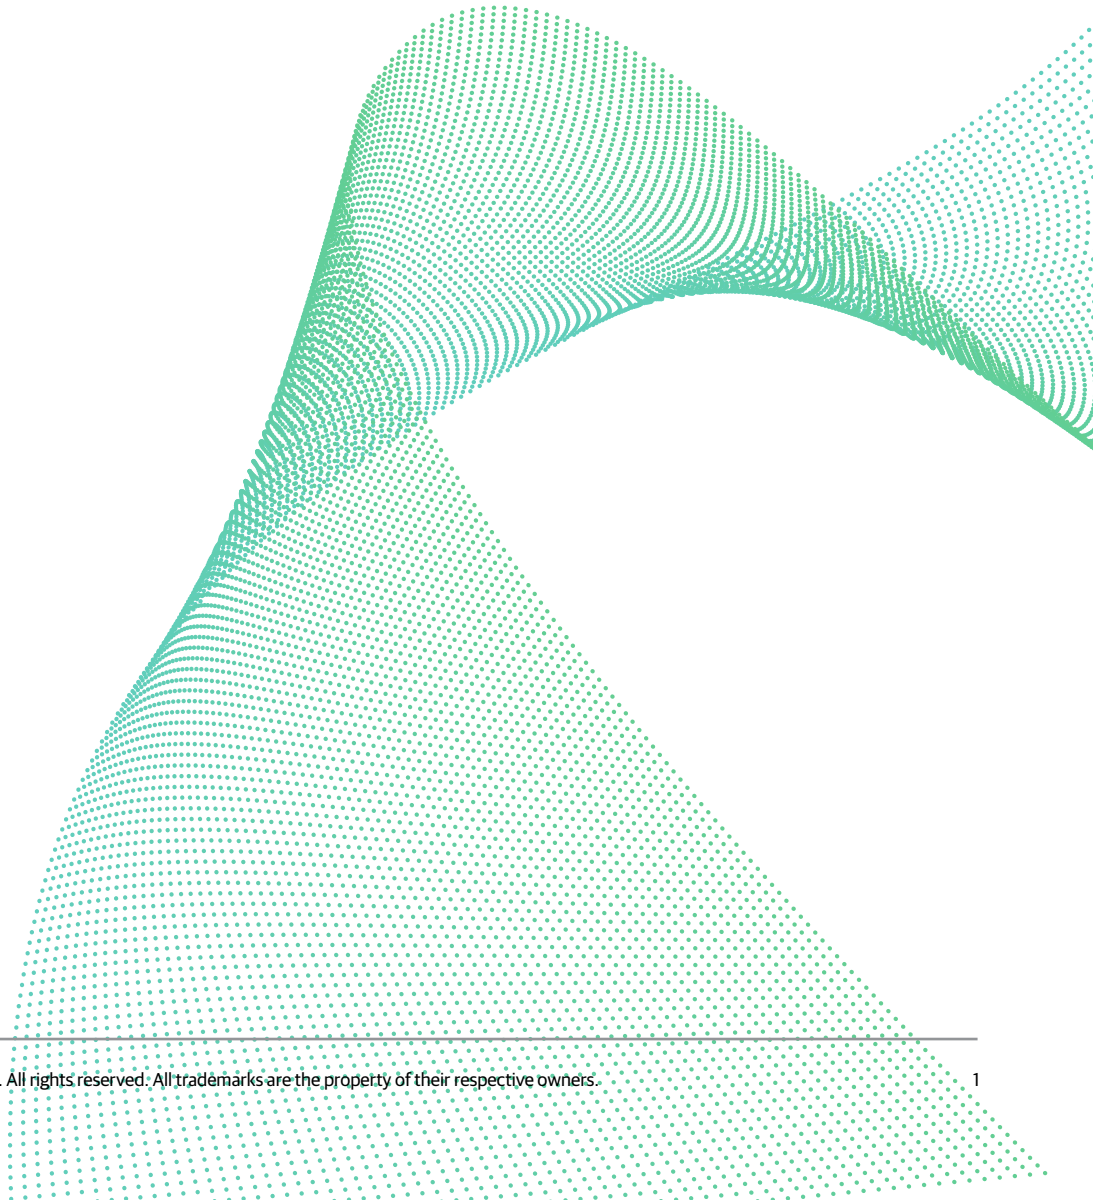

Contents

Disaster can strike anywhere at **anytime**, and often it impacts people that cannot afford to have a disaster. We are going to look at some things today that will help you more easily plan for a disaster-related outage, so that you can survive it.

Definitions

It is worth understanding some of the terms that are used in this discussion — starting with what exactly DR is — we know that it means disaster recovery, but what does that mean? Making Exchange work again is DR. No one can access it, and no email can be delivered to it, but it is running. BC is business continuity, and it means to bring back the ability to work, which would be accessing Exchange and sending and receiving email. In terms of BC, there is recovering Exchange, but also the internet and other things that would allow the sending and receiving of email.

You also often hear about RTO and RPO. RTO is recovery time objective, and it refers to how much time goes by before you can be back at work. RPO is recovery point objective and refers to how much data is lost before you are working again. The closer that either or both are to zero, the larger the cost is. Most applications and most customers do not need near-to-zero RPO and RTO due to the cost and the complexity. A real-life example of an organization requiring a near-zero RPO and RTO would be Amazon. The online purchasing giant would have a zero for RTO and RPO since the amount of money they would lose in 10 seconds of outage is measurable. However, accounts payable and accounts receivable would likely be in the two hours of RTO since there is no need to have zero for them. The cost difference between zero and two hours is likely breathtaking.

Applications

Something else that is important to consider is what is an application? Continuing with the Exchange example from above, it is important to understand that in DR work, Exchange is not an application. Exchange, Active Directory Domain Controller, DNS, DHCP and a desktop (for testing) are the application. By thinking this way, you can have easier testing and portability because you can support granular failover.

It is very useful to recover email first as it will allow you to communicate with outside parties to assure them that all is OK. This can be employees, the press or analysts, which really helps the market think you are better off than perhaps you are, and that is important.

Next, you should recover accounts payable. If you can pay your bills through the DR event, especially if it takes a while to fully recover, you will look better to the market, and it will help you deal with things such as market perceptions better.

Don't forget you do not need to protect everything. Things like monitoring tools should already be at the recovery site, including Activeincluding Active Directory, DNS and DHCP, so there is no need to protect them.

Sometimes it is hard to package an application as you do not know the pieces to it. In this case, you can talk to the help desk and see who is responsible for the most support calls or who is Tier-3 support, and eventually you will find who owns the software, but also who knows the most about it. After that, you have the information that makes packaging the app easy.

VMware has a tool called VMware Infrastructure Navigator (VIN) that will show you what VMs and processes talk to other VMs and processes. This is quite helpful when you are building your application packages. In addition, testing of an application during DR testing often fails due to missing components. This is something that VIN can help you avoid. VMware has another tool that might help with this as well called vRealize Operations Service Discovery Management Pack.

Practice

Practicing frequently is what makes the difference between succeeding and not. One idea is to practice a test failover (not an actual failover) of some application every Friday afternoon. It will likely not work on the first try, but you just need to keep at it. Application owners will start learning more about their application and how to do a successful failover.

Once the failovers start working all the time for an application, start sabotaging it. Old domain controllers are a good way to confound people!

Another idea I like is once the Exchange team is good at successful test failover (and troubleshooting), have them switch with the SharePoint team.

This is not about making people unhappy but about realistic practice, so when a true DR event occurs, the odds of successful failover are very high.

You do test failovers until everyone knows what is going on and is confident, and then you choose one single application and do a real failover — perhaps on a weekend — and see how it goes. With each successful test, you add an application until you are very confident.

Statistics

You will likely hit a DR event in your career — maybe even a few times. See below for some statistics on the effects of a DR event:

- 94% of companies suffering from catastrophic data loss do not survive, 43% never reopen and 51% close within two years (University of Texas). Gartner says something similar and so does Continuity Statistics.
- 18.5 hours is the length of an average outage. The cost for a small company to recover is $8,000 and is $700,000 for a large enterprise (IntraScale 2015).
- A five-hour computer outage costs $150 million (Delta — 9/7/16).
- In the Veeam® 2017 Availability Report, one of the findings was that four out of five organizations have an Availability Gap. This means that 82% of the respondents recognized the inadequacy of their recovery capabilities when compared with the expectations of the business unit.
- Also from the Veeam Availability Report was the fact that 72% of respondents are unable to protect their data frequently enough to ensure access to the data.

I share these statistics to help you understand that it is probable you will have a DR event, so you should prepare for it.

Good stuff to know

One of the best ways to reduce your outage is to have some of your apps running on the recovery site already. As part of your DR preparations, you should think about what can be permanently moved to the recovery site.

Applications that are geo balanced — meaning that they can handle outages without end users knowing and that, when you connect, you are connected to the closest copy of the resource — often do not need any DR. However, sometimes they might need DR in lower levels of the application stack, so be aware. Hopefully if you are lucky though, no DR will be required.

Always have an executive sponsor. DR work goes across great sections of your company and sometimes that makes things too difficult to be successful, so an executive sponsor is very handy.

Most DR events are partial, which is why you organize and protect for that. An example would be a building power outage. That is very worrisome and often you will migrate one or two applications to another site, so that when the building power outage of two hours **actually lasts** 24 hours, most of your users will not notice because you have email working for them in the recovery site.

As you get good at DR, your support people and application owners get much better at understanding their applications, and your company gets confident in the DR work. This is important as it means you can now do preventative migrations. A storm is coming, and it looks like a bad one, so you can move some of your apps inland to your recovery site. Even if the storm is not so bad, you have again practiced your DR and a short outage was all that was necessary to have some of your apps running elsewhere.

Start small by protecting something easy and small first. People will learn more easily that way, and things will get proven more easily too. With that experience, it is easier to do the bigger stuff.

Make sure to track the work you do — some sort of audit trail and reporting should be in every DR tool. Failovers are not that hard, but making the outage **really short** is tough. By keeping the reports and audit trails, it can help you with designing things better to minimize the outage.

IP addresses

Why do IP addresses get their own section? A typical DR tool can manage Windows and Linux Operating System IP addresses, so it's no big deal. Most of the DR recovery sites I have worked with have a different IP scheme, so this is important. But the real issue is that **sometimes homegrown** applications have IP addresses compiled in. This means that when they failover to the recovery site, the OS gets the IP address changed, but the app doesn't work. There are expensive and complex solutions to this, but it is best to avoid it if you can. Test carefully so you know in advance. Sometimes there are applications that need an IP address and they store it in a text file — tnsnames.ora is an example. Most DR tools can manage that sort of issue via a script that executes inside the VM and edits the text file.

Trigger time

There are two components to trigger time: the decision to execute your DR plan and the actual button push. Both are important and both need to be practiced. Agree in advance who can make the decision to execute and who can push the button. When things go badly and you need to execute your plan, you will not have the people that practiced the most, but rather the people that practiced the least, so it is very important the trigger process is well known. It is very important that the DR plan is triggered when it is needed, and we do not want the lack of knowledge to delay this decision to delay it!

My examples

I was a professional services fellow for a long time. I was lucky enough to be involved with a lot of interesting companies, and, **as a result**, I learned a lot about BC, DR and backup. Here are some examples of what I went through that helped me learn a lot about DR:

- A mini data center onboard a ship had a fire, which was something we were prepared for and had several plans in place to deal with — not the least of which was to power off fast and dry things out fast. We even had extra gear. I did not realize at the time, however, that sea water is what is used to fight fire, so drying out the gear did not work.
- I had several clients that needed to be connected to the internet **at all times**. I configured two connections via two different ISPs to protect them. Sometime later, they lost the internet connection because a train car went off the rails in the mountains and hit the side of a hill, where there was a pipe with both ISPs' fiber in it.

- I helped a customer move from water sprinklers in the data center to INERGEN. After that was completed, we had some water damage. It turns out that the workers terminated the water line just outside the data center, so when it leaked, it still impacted the data center.

Another example that I think is important but did not live through was an employee who heard that he was going to be let go. Before that happened, this employee **opened up** several hundred of the most important files, did a ctrl A, followed by ctrl Z and then ctrl S. So, a lot of important files are now empty.

These examples are of things that could have gone better, but we also had multiple methods of dealing with them. You need to pick what fits into your BC strategy. In the last case, you have the option of a backup restore or the option to do a DR failover, and then you can leave the evidence for the law and investigation to deal with.

Tools

The first tools were scripts, batch files and check lists. They worked. Customers who have experience with these first tools are always more successful with using other tools. Veeam has a new tool, and more information on it can be found here.

Why do you need tools? Generally, they make it easy and safe to practice. After all, you do not want practice data in production! Without tools, practice is tricky and generally expensive, which means it is not done very often.

Why does DR fail?

- Because BC is needed!
- Attitude is missing
- Executive sponsor is missing
- Missing application knowledge
- Failure to practice
- Data from test gets into production

Summary

- Know your applications
- Organize to support practice
- Practice like things are real
- Be ready!

I hope this white paper empowers you to be ready for whatever DR events you may experience.

About the Author

Michael White started out after leaving the military in Professional Services for a VMware partner. After doing technical implementations around the world, he joined VMware. Michael started as a Partner SE, followed by Specialist SE, then Staff Technical Marketing Architect, and finally into R&D as an Integration Architect. Much of his career has been in BCDR related work. Michael have spoken on a variety of topics at TSX, PEX, VMworld, Gartner, and local VMUGs. He started at Veeam in Technical Marketing specializing in the Veeam Availability Orchestrator product but recently moved into Research & Development.

About Veeam Software

Veeam® recognizes the new challenges companies across the globe face in enabling the Always-On Business™, a business that must operate 24.7.365. To address this, Veeam has pioneered a new market of Availability for the Always-On Enterprise™ by helping organizations meet recovery time and point objectives (RTPO™) of < 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified protection, leveraged data and complete visibility. Veeam Availability Suite™, which includes Veeam Backup & Replication™, leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs.

Founded in 2006, Veeam currently has 47,000 ProPartners and more than 242,000 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit http://www.veeam.com.

AVAILABILITY for the Always-On Enterprise™

VEEAM

Veeam makes the Fortune 500 Available. 24.7.365

To enable its **Digital Transformation,** 70% of the Fortune 500 rely on Veeam to ensure Availability of all data and applications. **24.7.365**