druva

Checklist

5 STEPS: AUTOMATE DISASTER RECOVERY WITH THE CLOUD

Q317-CON-10699

Introduction

While the rise of virtualization has dramatically enhanced server and primary storage utilization, the threat to your virtualized environment is now more acute than ever. Prudent organizations already have a comprehensive on-demand solution in place to provide immediate data spin-up and failover to the cloud should a disaster occur. Snapshots of virtual machines stored in the cloud enable enterprises to achieve multiple data protection workloads within a single environment. Cloud-native solutions offer a zero-infrastructure, scalable solution with dramatically lower total cost of ownership (TCO), while delivering all the benefits of offsite data protection and recovery.

Checklist

The following checklist describes five simple steps that you can—and should—take today to set up and automate a single-click DRaaS solution for your organization.

1. Set up Your Storage Regions

This first critical step will be largely determined by the location of your remote offices and by various regulatory requirements. For example, a company with offices dispersed worldwide may want to ensure that data is backed up in the same region where the office resides, while sensitive PII data (such as data that's covered by HIPAA) may be limited by statute to even more strict geographic boundaries. This helps make data spin-up in the event of a disaster more efficient and allows for a recovery time objective (RTO) of mere minutes. Decide where your data needs to be stored, and ensure that your chosen vendor provides the coverage that your organization requires.

2. Prepare Your Network

Once you have decided upon a location for your disaster recovery (DR) data, it is important to configure your network to access these stores. One issue that system administrators often have to deal with is inconsistency in network availability and connectivity in remote and branch offices. Define the location of the data to be backed up as well as the location of the backup server, making special note of the network topology and any bandwidth constraints. Get the network team involved at the beginning of this process, and define the key stakeholders that you'll need before you start testing and configuring the network based on expected traffic. As with any other project, working on solutions early to address issues that you can see coming is always the best strategy to ensure that your team, your infrastructure, and your budget are taking everything into account.

3. Define Your Backup Policies

Which data is most important for you to back up? All server data and VMs? End-user data? And what is the required frequency for snapshots? All of these policies will need

to be defined based on your company's unique situation. In many cases, a recovery point objective (RPO) of no more than a couple of hours may be required—possibly even less for key data and executives. These policies will be defined and managed from within the console of your DRaaS solution.

4. Register and Configure Your Servers

To ensure that your virtual machines are being backed up successfully to the locations you have predetermined, it is critical that you register your servers with the cloud DR service. To ensure proper security, an effective DRaaS product will require that each server be authorized. Once you have defined and authorized your servers, you can configure VMs by assigning them to specific server groups for backups. You will also need to define firewall rules that will allow communication between the client application and the backup server. In the event of a disaster, there's nothing worse than discovering that your network traffic has been blocked by an overzealous but well-meaning security suite.

5. Configure Your DRaaS

If a disaster occurs, you need to be able to recover data quickly and with minimal impact to your employees and your business operations. Ensure that you have all critical virtual images stored appropriately within your DRaaS management console so that they are ready for instant failover to their respective sites. In addition, define which users will be tasked with administering the deployment. Would it be better to have a credentialed administrator at each site, or will one person or office be responsible for the entire system? Also, consider using one-to-many replication, which would allow your organization to manage VMs across geographies and accounts for cloud migration and dev-test purposes with no additional infrastructure needed.

Conclusion

Unfortunately, getting hit with a disaster event like ransomware has become a case of when, not if. Storing any significant amount of business data without a DRaaS solution in place is akin to a game of Russian roulette: the longer you play, the greater your odds of calamity become, until it is eventually all but certain. By following the steps in this checklist, you can set up and automate a comprehensive DR deployment to allow for single-click failover of your mission-critical data and have your organization up and running again within minutes.

Learn how to enable DR automation for your organization at
https://www.druva.com/solutions/cloud-disaster-recovery/

About Druva

Druva is the leader in cloud data protection and information management, leveraging the public cloud to offer a single pane of glass to protect, preserve and discover information—dramatically increasing the availability and visibility of business critical information, while reducing the risk, cost and complexity of managing and protecting it.

Druva's award-winning solutions intelligently collect data, and unify backup, disaster recovery, archival and governance capabilities onto a single, optimized data set. As the industry's fastest growing data protection provider, Druva is trusted by over 4,000 global organizations and protects over 25 PB of data. Learn more at http://www.druva.com and join the conversation at twitter.com/druvainc.

Druva, Inc.
Americas: +1 888-248-4976
Europe: +44.(0)20.3150.1722
APJ: +919886120215
sales@druva.com
www.druva.com