SearchITChannel

TechTarget

E-Guide

Cloud Disaster recovery Guide

Contents

***No one lives in a protective bubble. So as long as** there are natural disasters and technical failures, customers will continue to regard disaster recovery (DR) and cloud-based Disaster Recovery as a Service (DRaaS) as valuable services. This eGuide will you brush up on the latest DR tips, definitions and other resources to successfully enter the DRaaS provider market and remain competitive.*

Planning cloud disaster recovery services and avoiding the pain points

Disaster Recovery as a Service (DRaaS) is an up-and-coming service that can nicely round out a cloud provider's product portfolio. DRaaS is a low-cost alternative to expensive and often unwieldy traditional disaster recovery options. But cloud disaster recovery services are more than just cloud storage; they encompass planning, process, integration, testing and constant vigilance. Done wrong, DRaaS can spell disaster for your customer's business and your reputation. Here's how to avoid the pitfalls and successfully enter the DRaaS market.

Who knew what a derecho was until this unusually violent type of storm walloped Northern Virginia and resulted in a serious Amazon Web Services outage for its Elastic Compute Cloud 2 (EC2)? The event brought down many of Amazon's big-name customers, including NetFlix, Dropbox, Pinterest, Instagram and Heroku, along with a long list of other unfortunates. The message is that any enterprise whose business interests would be seriously harmed by a service disruption is a potential candidate for DRaaS.

Who wants cloud-based disaster recovery services and why?

Traditional disaster recovery services replicate application state between two data centers, so if the primary data center becomes unavailable, the backup site can take over and activate a new copy of the application using the most recently replicated data. Disaster recovery services are well suited for the cloud because under normal conditions, minimal resources are needed to synchronize state from the primary site to the cloud. The full complement of resources required to actually run the application need only be provisioned and paid for when disaster strikes. Additional resources can be rapidly called upon for quick recovery after a failure, enabling business continuity.

Now, cloud-based DRaaS is a viable alternative to traditional disaster recovery if a customer's applications are already running in a virtualized environment. Applications running on mainframes and hybrid computing environments are best served by traditional disaster recovery offerings.

SunGard, which has been offering traditional data disaster recovery services for more than 30 years, added cloud-based disaster recovery last year. According to Ram Shanmugam, senior director of product management for recovery services at SunGard Availability Services, about 20% of SunGard's customers run their applications in fully virtualized environments -- and that segment now uses SunGard's DRaaS offering. Shanmugam predicts that the percentage of customers with fully virtualized environments will grow anywhere from 20% to 60% over the next two years. If his prediction is accurate, DRaaS will quickly come into its own.

How to implement DRaaS and avoid the pain points

Implementing DRaaS solutions entails creating a cloud disaster recovery plan with Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs) based on business considerations and well-defined processes for recognizing and declaring a disaster. It also entails tight integration between the network, firewalls and load balancers, and the Web, database and storage tiers. Storage is at the bottom of the stack, and, according to Shanmugam, DRaaS must be equipped to recover the entire stack above storage or the recovery will fail.

A thorny DRaaS implementation issue is that all operating system software for the original site and the backup site must be monolithic in terms of vendor and version number. This makes it challenging to provide DRaaS offerings from a customer's private cloud to your cloud or between service provider clouds. Any time software changes, incompatibilities can be introduced and things can break.

Another implementation problem is that vendor solutions to date have been designed for enterprises and lack a feature set upon which to build cloud-based, multi-tenant solutions. As a result, cloud providers have been forced to do their own software development.

Craig McLellan, CTO of Hosting.com explains the problem. "There is a school of hard knocks involved [in implementing cloud Disaster Recovery as a Service]. One of the key things we've learned is that virtually none of the vendors who sell their wares to service providers knows anything about offering disaster recovery services on a multi-tenant cloud footprint. The APIs available from vendors are woefully inadequate. They don't expose the hypervisor to the customer, so we have to write the software to provide access to the customers to test and administer the service. None of the vendors contemplated that. If they had an API, we could create multi-tenant abstractions -- without that, it is very difficult."

McLellan also points out that vendors are designing products for what disaster recovery used to be. In the past, testing happened annually; RPOs were not as aggressive. "Now we need to test more often, and to do that, the testing must be 100% nonintrusive. Currently we have to do it manually, and that is intrusive. What is needed is a reliable orchestration layer in the middle," he said.

Good performance data isn't readily available just yet, McLellan continued. For a service provider, service-level agreements (SLAs) are key -- therefore performance is key. "Without good instrumentation, we can't offer SLAs. God help you if you can't tell how your service is performing and a customer declares a disaster. It can easily destroy a service provider's reputation."

He said his company overcompensates by putting in its own instrumentation or by putting a person on it because service providers don't get a second chance. "There is no forgiveness in this space. There's lots of competition, and the switching cost is relatively low," he said.

Partnering for more effective DRaaS offerings

So if you plan to jump into the DRaaS market with a do-it-yourself solution today, be prepared to blaze your own trail. If tackling that option may be low in your priority queue, however, you may be able to enlist a partner to help jumpstart your DRaaS market entry.

Given that it is not easy to roll out your own DRaaS offering, you may want to consider outsourcing to a third party. SunGard resells its DRaaS to cloud service providers under the SunGard brand, which has caché in the disaster recovery space. But you can white label the service if you choose.

You can also work with a cloud system integrator like CloudOps. According to Ian Rae, CEO at CloudOps, "We can architect a cloud-based disaster recovery solution and manage it on behalf of a customer. The advantage of working with a cloud system integrator like us is that we have a strong understanding of what is involved in restoring an application to service after an outage."

The future of cloud disaster recovery services

Cloud disaster recovery services are likely to catch on quickly, and forward-thinking cloud providers should keep their eyes on the horizon to take advantage of what might come next. One possibility is service provider diversity. Just as geographic diversity is important for disaster recovery planning, service provider diversity is also important. In preparation, look for ways to provide DRaaS to customers of other cloud providers, which would help you expand your customer and revenue base and possibly win over additional service business in the long run.

DRaaS definitions you need to know

Cloud service providers thinking about adding Disaster Recovery as a Service (DRaaS) need to add terms like Recovery Point Objective, Recovery Time Objective, hot, warm and cold backup sites, and geographic diversity to their vocabularies. Here's a list of the essential terms: Recovery point objective. An RPO is the targeted maximum amount of time that can be tolerated between mirroring your data. Some enterprises may require an RPO of zero, meaning they cannot lose *any* data. This requires continuous, synchronous replication. Other enterprises may be able to tolerate data gaps of seconds, minutes, hours or even days if they have to revert to a secondary site.

- **Recovery time objective.** An RTO is the target interval between when an application outage occurs and when the application must be back up and running. This includes the time it takes to detect the failure, prepare the backup site, initialize the failed application and perform any network configuration to reroute requests to the backup site. The lower the RTO, the shorter the time between disaster and recovery.
- **Hot, warm and cold backup:** A hot backup site typically houses mirrored standby servers that are always available to run an application after a disaster. This is an expensive option favored by enterprises with short RTOs and RPOs. A warm backup site generally keeps state up to date, but it takes additional time to bring online because the resources are not immediately available. A cold backup site often has an RPO of hours or days, and data is often replicated in increments of hours or days. It is a low-cost option for applications that do not require rigorous availability guarantees.
- **Geographic diversity.** Another key disaster recovery term, geographic diversity, is important because primary and backup sites should be distant enough from each other to minimize the possibility that a single disaster takes down both sites.

Which disaster recovery solution for the cloud?

According to a report by researchers at the University of Massachusetts and AT&T Labs, warm standby services are best suited for cloud disaster recovery services. "Cloud platforms can provide the greatest benefit to DR services that require warm standby replications," the report said. "In this case, the cloud can be used to cheaply maintain the state of an application using low cost resources under ordinary operating conditions. Only after a disaster occurs must a cloud-based DR service pay for the more powerful --

and expensive -- resources required to run the full application, and it can add these resources in a matter of seconds or minutes."

Contents

Free resources for technology professionals

TechTarget publishes targeted technology media that address your need for information and resources for researching products, developing strategy and making cost-effective purchase decisions. Our network of technology-specific Web sites gives you access to industry experts, independent content and analysis and the Web's largest library of vendor-provided white papers, webcasts, podcasts, videos, virtual trade shows, research reports and more —drawing on the rich R&D resources of technology providers to address market trends, challenges and solutions. Our live events and virtual seminars give you access to vendor neutral, expert commentary and advice on the issues and challenges you face daily. Our social community IT Knowledge Exchange allows you to share real world information in real time with peers and experts.

What makes TechTarget unique?

TechTarget is squarely focused on the enterprise IT space. Our team of editors and network of industry experts provide the richest, most relevant content to IT professionals and management. We leverage the immediacy of the Web, the networking and face-to-face opportunities of events and virtual events, and the ability to interact with peers—all to create compelling and actionable information for enterprise IT professionals across all industries and markets.