



Denny Cherry
Denny Cherry
& Associates
Consulting

STOP QUAKING IN YOUR BOOTS ABOUT DISASTER RECOVERY

It will be a while before Southern California has all the data on production dollars lost from the Mojave earthquake, but it's a virtual guarantee that it will be painful for some small or enterprise businesses. As a certified IT expert and published author in data security and cloud technologies, I meet regularly with CEO's and CIO's who are aware their Disaster Recovery strategy has significant flaws, yet they don't budget to solve the problem. The reason is usually because their DR strategy is following an outmoded philosophy and structure as an earnings drain: space in a data centre, sitting idle and empty, racking up costs, waiting for production to have a critical failure.

If this is you, it's time to think about your DR strategy as a necessity for business continuity, rather than recovery. It's also time to think about how you can optimize that strategy to improve your business. There is an often-overlooked method to making your DR strategy cost-effective while enjoying a return on your investment: by diversifying production so the DR site is utilized to improve customer experience and generate revenue.

Here's how it works. Setting the stage for a moment, let's talk about a sample company with its production data centre in California and its Disaster Recovery data centre in New Jersey. The company employees are in California; while the customers are across the United States. With a traditional DR configuration, the customers would all connect to the site in California, and if there was a disaster at that facility the customers would all be routed to the site in New Jersey.

With some redevelopment of the applications that are used by the customers, the application can be hosted in both the New Jersey and California data centres. Users of the application would be routed to the copy of the application, which is geographically closer to them, increasing the responsiveness of the application, allowing the company to potentially process revenue more rapidly, and in any case, increasing customer satisfaction.

For websites and web-based applications, this is a very easy configuration to make. The files behind the website are simply copied to web servers at both sites. A global load balancer,



which would be necessary for a proper disaster recovery environment, is configured to route users to one copy of the website or the other depending on where the user is connecting from. If a global load balancer hasn't already been purchased, using a service like Azure Traffic Manager or AWS Global Accelerator can do the same thing without having to put any other components within the cloud.

The most complex part of the application to configure to run in two (or more) sites is the database. By their nature, databases are single write, as concurrency needs to be maintained. However, in most applications, up to 95% of access to the database are reads, leaving only 5% of queries changing data. All big database vendors offer a way to read data from a secondary node, so it doesn't matter which database server is being used, the reads can be

done from a copy that is stored in the local site, and only the writes from the secondary site need to be done at the primary site. In the event of a failure of the primary site, only the writeable copy of the database needs to be failed over to the second site, as all the other services are already running there.

If performance of the application in the New Jersey site in our example isn't as fast as it should be, then Redis caching can be used for most queries. This caching layer removes load from the database servers by increasing the amount of data being stored and accessed in a caching tier, instead of having to go across the network for every database call. Not every database query can be cached, but most database calls can be, which greatly reduces the number of queries which are executed against the database.

When all of this is configured correctly, the application is a geographically distributed application stored in two (or more data centres) with all copies in all the data centres being used to generate revenue. This sort of configuration reduces the expenses for unused hardware to zero, as all the services are running from all the data centres and servicing customer requests. If configurations need to be scaled even further, they can

be. This same approach can be used to host the application in more data centres as needed, depending on where users are located.

Configured this way, Disaster Recovery is a benefit. One that results in increased uptime, higher customer satisfaction, and a higher utilization of the servers within your DR site, while guaranteeing business continuity... even if your wedding china doesn't fare quite as well in the quake.



Denny Cherry
& Associates Consulting

ABOUT THE AUTHOR

Denny Cherry is the author of "The Basics of Digital Privacy" and a world-renowned IT expert, speaker in data security and cloud technologies. A Microsoft MVP, he is CEO of Denny Cherry & Associates Consulting, a Microsoft Partner company.