

Disaster Recovery: Not Just for Data Centers Anymore



Lori MacVittie, 2010-10-11

It's about business continuity between the customer or user and your applications, and you only have control over half that equation.



Back in the day (when they still let me write code) I was contracted to a global transportation firm where we had just completed the very first implementation of an Internet-enabled tracking system. We had five whole pilot customers and it was, to say the least, a somewhat fragile system. We were all just learning back then, after all, and if you think integration today is difficult and fraught with disaster, try adding in some CICS and some EDI-delivered data. Yeah – exactly.

In any case it was Saturday and I was on call and of course the phone rang. A customer couldn't access the application so into the car I went and off to the office to see what was going on (remote access then wasn't nearly as commonplace as it is now and yes, I drove uphill, both ways, to get there).

The thing was, the application was working fine. No errors, nothing in the logs; I was able to log in and cruise around without any problems whatsoever. So I pull out my network fu and what do I find? There's a router down around Chicago.

Bleh. Trying to explain *that* to the business owner was more painful than giving birth and trust me, I've done both so I would know.

Through nearly a decade of use and abuse, the Internet itself has grown a lot more resilient and outages of core routers rarely, if ever, happen. But localized service (and its infrastructure) *do* experience outages and interruptions and there are just as many (if not more) folks out there who don't understand the difference between "Internet problem" and "application problem." And you won't have any better luck explaining to *them* than I did a decade ago.

GLOBAL APPLICATION DELIVERY to the RESCUE (SOMETIMES)

When most people hear the words "disaster" and "recovery" together their eyes either glaze over and they try to run from the room or they get very animated and start talking about all the backup-plans they have for when the alien space-craft lands atop their very own data center, complete with industrial strength tin-foil hat.

Okay, okay. There are people for whom disaster recovery is an imperative and very important, it's just that we hope they never get to execute their plans. The thing is that their plans have a broader and more generalized use than just in the event of a disaster and in fact they should probably be active all the time and addressing what external "disasters" may occur on a daily basis.

Global application delivery (which grew out of [global server load balancing](#) (GSLB)) is generally the primary strategic component in any disaster recovery plan, especially when those plans involve multiple data centers. Now, global application delivery is also often leveraged as a means to provide [GeoLocation-based load balancing](#) across multiple application deployments as a way to [improve performance and better distribute load across a global web application presence](#). If you expand that out just a bit, and maybe throw in some cloud, you can see that in a disaster recovery architecture, global application delivery is pretty key. It's the decision maker, the router, the "thing" on the network that decides whether you are directed to site A or site B or site Z.



Because many outages today are often localized or regionalized, it's often the case that a segment of your user population can't access your applications/sites even though most others can. It's possible then to leverage either regional clouds + global application delivery or multiple data centers + global application delivery to "route around" a localized outage and ensure availability for all users (mostly).

I keep saying "mostly" and "sometimes" because no system is 100%, and there's always the chance that someone is going to be left out in the cold, no matter what.

The thing is that a router outage somewhere in Chicago doesn't necessarily impact users trying to access your main data center from California or Washington. It does, however, impact those of us in the midwest for whom Chicago is the onramp to the Internet. If only the route between Chicago and the west coast is impacted, it's quite possible (and likely) we can access sites hosted on the east coast. But if you only have one data center – on the west coast – well, then you are experiencing an outage, whether you realize it or not.

cloud computing MAKES MULTIPLE DATA CENTERS FEASIBLE

For many organizations dual-data centers was never a possibility because, well, they're downright expensive – both to buy, build, and maintain. Cloud computing offers a feasible alternative (for at least some applications) and provides a potential opportunity to improve availability in the event of *any* disaster or interruption, whether it happens to your data center or five states away.

You don't have control over the routes that make up the Internet and deliver users (and customers) to your applications, but you do have control over where you might deploy those applications and how you respond to an interruption in connectivity across the globe. [Global application delivery](#) enables a multi-site implementation that can be leveraged for more than just disaster recovery; such implementations can be used for localized service disruption as well as overflow ([cloud bursting style](#)) handling for seasonal or event-based spikes or to assist in maintaining service in the face of a DDoS attack. "Disaster recovery" isn't just about the data center, anymore, it's about the applications, the users, and maintaining connectivity between them regardless of what might cause a disruption – and where.



Related blogs & articles:

- [Location-Aware Load Balancing](#)
- [Cloud Needs Context-Aware Provisioning](#)
- [What is a Strategic Point of Control Anyway?](#)
- [So You Put an Application in the Cloud. Now what?](#)
- [The Inevitable Eventual Consistency of Cloud Computing](#)
- [Cloud Needs Context-Aware Provisioning](#)
- [Cloud Balancing, Reverse Cloud Bursting, and Staying PCI-Compliant](#)
- [Pursuit of Intercloud is Practical not Premature](#)

- [Cloud Balancing, Cloud Bursting, and Intercloud](#)
- [Getting Around That Pesky Speed of Light Limitation](#)

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

©2016 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. CS04-00015 0113