

Enterprise Networks



STEVE GUENDERT, PH.D.

Control Your Data's Destiny

Earlier this year, while on a business trip to Europe, a client shared an interesting story with me when we were discussing high availability and network architectures. I'm going to share that story, but in the tradition of the old TV show "Adam-12," I've "changed the names to protect the innocent." The story goes something like this:

This company had decided to build a new data center approximately 15 km away from their production data center, which was in a suburb. They already had another data center that was approximately 125 km away from the production data center, and they were doing asynchronous remote DASD copy between it and the production data center. They were going to implement synchronous remote copy between the production data center and the new data center. In other words, they were moving to a typical, three-site Disaster Recovery (DR)/business continuity architecture.

The synchronous copy would be done over their extended FICON SAN between data centers, and the network design requirements called for a high-availability/multipath design. At this point, the mainframe and mainframe storage teams' involvement in the network architecture and design process stopped and was turned over to "the network guys," who told the mainframers they understood the network requirements and would implement them. The network guys met with the two service providers (they contracted with two because that "meant it would be redundant paths") and came back to the mainframers with a network design diagram. The network design diagram showed one path between data centers going around a ring on one side of the city and the other path going around the other side of the city with a different service provider.

Everything was built, put into production and all worked fine for several months. Then one day things went to pieces. Both network paths between the two local data centers went down nearly simultaneously. Not good. How could this have happened? What were the odds of this happening when you have redundant network paths going around the city, each with a different network service provider?

Well, actually the odds were pretty good. You see, as it turns out, the nice little network diagram didn't show the whole picture. There was a set of railroad tracks that went across the city. One of the service providers had its path in the right of way on the north side of the tracks and the

other had its path in the right of way on the south side of the tracks. The railroad decided to do some maintenance on its tracks, and the inexperienced, overzealous operator of the railroad's fibre optic cable hunter (aka backhoe) got a little carried away and took out both paths. To quote that wise old sage Homer Simpson, "Doh!"

What were the odds of this happening when you have redundant network paths going around the city, each with a different network service provider?

The lesson to take away from this unfortunate example is that as mainframers, we have a different idea of high availability and redundancy than most other people in IT. Our paradigm is based on what we've come to expect from our computing platform. We can't afford to assume that others (such as the network guys) will understand this. When it comes to network decisions, we need to:

- Clearly communicate in great detail what our requirements are and what we mean by high-availability redundant paths
- Not settle for being told "just tell us how many ports you need and we've got it from there"
- Demand a seat at the table for the details of the network decisions that potentially impact the data we're responsible for.

Better yet, maybe it's high time we tell the network guys we want our own dedicated DR network; we'll design it, purchase it and manage it. **ETJ**

Dr. Steve Guendert is a principal engineer and global solutions architect for Brocade Communications, where he leads the mainframe-related business efforts. A senior member of both the Institute of Electrical and Electronics Engineers (IEEE) and the Association for Computing Machinery (ACM), he serves on the board of directors for the Computer Measurement Group (CMG). He's also a former member of the SHARE Board of Directors.
Email: stephen.guendert@brocade.com
Twitter: @BRCD_DrSteve