

**STEVE
GUENDERT,
PH.D.**



A line of demarcation is defined as a line defining the boundary of a buffer zone or area of limitation. A line of demarcation may also be used to define the forward limits of disputing or belligerent forces after each phase of disengagement or withdrawal has been completed. The term is commonly used to denote a tempo-

rary geopolitical border, often agreed upon as part of an armistice or ceasefire. The most famous line of demarcation in recent history is the Military Demarcation Line, also known as the Armistice Line, which forms the border between North Korea and South Korea. There also was the late Libyan leader Muammar Gaddafi's ironically named "Line of Death." But enough history; you get the point. By now you're probably asking yourself how this relates to your data center.

I meet with many mainframe IT organizations across the globe, and one thing the vast majority have in common is this: There are lines of demarcation that exist to separate the responsibility for the various teams that manage the mainframe, FICON Storage Area Network (SAN), mainframe storage and the network. If you're running Linux on System z, there are likely more, but that's another story. Here I will focus on the line of demarcation that exists in many business continuity architectures; that's the line of demarcation between the team(s) that manage the mainframe/mainframe storage/FICON directors-channel extension and the team that manages the network for cross-site connectivity. I like to call this the Storage-Network Line of Demarcation.

The majority of you have two sites and a cascaded, multi-fabric FICON architecture connecting these sites. The distance between sites and your Recovery Point Objective (RPO)/Recovery Time Objective (RTO) will dictate if you're doing synchronous or asynchronous DASD replication across the network between sites. Many of you also have a form of mainframe virtual tape solution performing replication between sites via your network. Some of you have more than two sites. For example, you may have two sites located fairly close together and you perform synchronous DASD replication between those two sites. You then have another site located a long distance apart, and perform asynchronous DASD and/or virtual tape replication between them. If you're running IBM's Geographically Dispersed Parallel Sysplex, you have even more host-related connectivity traversing the network.

Question for my mainframe and storage friends:

Who's responsible for the network connectivity/hardware for your data traffic between sites?

In many of the organizations I visit, a line of demarcation exists between

mainframe/storage and network; for example, a client running XRC between sites and using a Fibre Channel over IP (FCIP) extension port card in their FICON director. That FCIP extension device will be connected to a network of some sort, which then connects to another FCIP device and FICON director at the remote site. More often than not, the team responsible for the FCIP devices, and the data traffic running between them, have no say when it comes to the IP network to which they're connected.

Guess where the problems usually are (especially if Dense Wave Division Multiplexing [DWDM] is involved)? Guess who typically is held responsible/gets blamed for the inter-site data traffic performance problems (hint: not the network team and their equipment)?

Or worse yet, in addition to the line of demarcation, I often see the network equipment, whether it be Internet Protocol (IP), DWDM, etc., that's old, slow, poor performing hardware that simply can't keep up with the newer FCIP and FICON director performance, hence creating a bottleneck. What does this do to your RPO and RTO?

Finally, I've heard that working with others on the other side of this line of demarcation is often like dealing with the fictional USA Prime Credit ("USA Prime Credit, this is Peggy") customer service team from TV commercials.

There's a better, less political way for you to have control of all the hardware in these networks; there's less infighting, with a simpler network with better performance that's easier to manage and an improved disaster recovery/business continuity posture for your enterprise. I think your CIO and CTO will like it, too. Consider a smaller, dedicated network for your disaster recovery/business continuity needs that you control. **EE**

The Storage-Network Line of Demarcation Correlation

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 is also a former member of the SHARE Board of Directors. Email: stephen.guendert@brocade.com Twitter: @BRCD_DrSteve