Improving Efficiency in Your Long-Distance Networks

STEVE GUENDERT, PH.D.

The vast majority of *Enterprise Executive* readers work for an organization that has multiple data centers. These data centers typically are geographically separated by a distance deemed by risk analysis to be safe from being impacted by the same disaster. If a regional disaster is a threat, the trend is toward a business continuity architecture of three sites or maybe more. Two of these sites will likely be in closer proximity to each other (10 to 25 km) with synchronous DASD replication between the sites. The third site will likely be much further away, with asynchronous DASD replication between it, and one or both of the other sites.

These types of business continuity architectures are expensive, but compared with the financial and/or business reputational costs of a major outage, they're well worth it. So, what do you think is the most expensive cost component of these business continuity architectures? The mainframe? The DASD arrays? The network hardware? If you guessed any of those, I'm sorry but you're likely wrong. I can't comment on your specific costs, but based on my experience with customers and clients worldwide, I can safely say that as a rule, the most expensive cost component of business continuity architectures is the network bandwidth between the data centers.

Here I will provide some ideas on how you can improve the efficiency with how you use that bandwidth. These ideas should help you reduce your bandwidth costs. A subsequent article will expand on these ideas in more detail.

First, you should always do a bandwidth needs analysis/assessment. This can be done in-house using commercially available software, or you can engage one of your vendors. For example, if the predominant bandwidth requirement is driven by XRC, you may want to engage your DASD vendor. A good bandwidth needs assessment will include a forecast for growth because, let's face it, your data is growing, and with it your replication requirements, hence, bandwidth requirements.

Next, take advantage of protocol efficiencies with newer technologies. For example, if your architecture includes cascaded FICON directors with Interswitch Links (ISLs) between sites, the latest FICON switching is 16 Gbps. Today, the primary use case for this is with ISLs. The reason is that 16 Gbps FICON uses 64b/66b encoding, while 8 Gbps and 4 Gbps use 8b/10b encoding. The overhead of the 64b/66b encoding is 3.125 percent, which is considerably less than the 8b/10b encoding scheme, which has a 25 percent overhead. Why use that expensive bandwidth for overhead when a much more efficient option is available?

Third, keep data streaming over the distance. Try to avoid bursty data traffic patterns. Going back to our previous cascaded FICON example, make sure you have plenty of buffer credits configured on the director ISL ports. It's also a good idea to use your SAN vendor's buffer credit recovery mechanism. Buffer credit recovery is a standards-based technique that detects and remedies buffer credit loss issues caused by things such as faulty Small Form-Factor Pluggables (SFPs).

> So, what do you think is the most expensive cost component of these business continuity architectures? The mainframe? The DASD arrays? The network hardware?

Finally, give serious consideration to using some form of trunking, as well as Quality of Service (QoS). Trunking is used to aggregate individual cross-site links into one or more "big pipes." QoS is a feature/management technique that allows the end user to prioritize specific workloads. This prioritization is extremely useful in the event one or more cross-site network links goes down and there may not be enough bandwidth left on the remaining links to maintain performance. QoS features allow you to ensure that in this event, the most important workloads have use of the bandwidth.

Those are some basic ideas for gaining more efficiency in your long-distance network. More efficiency means lower bandwidth requirements, which can significantly lower costs in your business continuity architecture. **EE**

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