Pipeline
Technology for Service Providers.
www.pipelinepub.com Volume 15, Issue 6

How Financial Services Companies Can Rest Assured of Their Network's Resiliency

By: Roy Hilliard

Network downtime isn't so much a question of if, but when, and the financial costs to businesses can be overwhelming, to say nothing of the adverse effects on productivity and brand reputation. According to a recent report by IHS Inc. (NYSE: IHS), network outages experienced by North American organizations alone result in revenue losses of $700 billion per year.

Even if a company's applications, servers and devices remain operational, when the network is down they no longer can communicate with one another. Moreover, these failures are not rare and isolated events. Research by Dunn & Bradstreet claims that 59 percent of Fortune 500 companies experience a minimum of 1.6 hours of downtime per week. The study by IHS found that organizations experience 27 hours of downtime per month.

Even a conservative estimate from Gartner projects the hourly cost of network downtime is $42,000, meaning that a company that experiences worse than average downtime of 175 hours a year can potentially lose more than $7 million annually. Meanwhile, an article in the *Financial Times* finds that average downtime costs vary widely across industries, from approximately $90,000 per hour in the media sector to nearly $6.5 million per hour for large online brokerages.

Needless to say, the potential of a $6.5 million hit—or even a much lesser figure—would keep any CIO up at night, if he or she had reason to doubt the resiliency of the company's network.

Network Downtime Does Not Discriminate

In the financial services industry, companies and consumers are more reliant on technology than ever before, making any business-critical IT outage all the more significant. Downtime can take many forms and repercussions, from an outage on a transatlantic, high-frequency trading platform to an email provider experiencing a service interruption, meaning customers can't receive their bank statements on time. Equally damaging, it could be a network outage that prevents customers from being able to take their money out of an ATM or perform a credit card transaction. Separate consumers from their purchasing power for even a few hours and you immediately see how it affects their customer loyalty.

These outage events are increasingly in the public eye, and despite whatever solace network

engineers may take in the promise of five-nines availability or low latency infrastructure, outages do not discriminate based on a financial institution's size or geographic location. In the US, the online and mobile banking platforms of Atlanta-based SunTrust's eleven branches went down due to a technical difficulty associated with a network upgrade. The major global banks PNC Financial Services, JPMorgan Chase and TD Bank have also experienced network outages that affected their retail and commercial customers. North of the border, a network glitch at CIBC, one of the 'Big Five' banks in Canada, caused interruptions to its online and mobile banking systems. Across the pond, UK-based financial institutions Lloyds Banking Group, Barclays and the Royal Bank of Scotland have all fallen victim to multiple network outages.

Most notably, all of the network failures listed above have occurred in only the previous twelve months. Even with five-nines, there remains a 0.001 percent opportunity for system failure.

The Case for Network Infrastructure Clarity

Although banks and other financial services institutions have become more forthright in issuing post-outage apologies on their social media platforms, they are rarely transparent about the actual cause.

Statements such as "We are currently experiencing a technology issue; our teams are working to restore full access as soon as possible," offer little insight and lesser comfort to customers and partners. But whether caused by hardware failure, network configuration issues, human error, routing issues and congestion, or catastrophic weather events such Hurricane Sandy, having a recovery plan—including rerouting traffic over a redundant link—is the bulwark against extended outages and financial loss.

To relate a story from the trenches, my colleagues and I once met with representatives from a global bank that deploys multiple MPLS networks in the US, a meeting that laid bare just how vulnerable New York is to system failures. Adopting a traditional approach, the financial institution purchases network connectivity from two carriers that unsurprisingly, but unknown to them, overlap in at least one or two nodes. When one of the organization's Points of Presence (PoPs) went down, applications went offline for several hours, which led to many uncomfortable conversations among the leadership team.

"How did this happen? More importantly, how can we prevent this from ever happening again?"

We should note that, as recently as five to ten years ago, a two-carrier strategy was adequate. But in today's connectivity landscape, especially in view of the recent surge of M&As and the increasing complexity of networks, it's becoming increasingly difficult to achieve a high degree of confidence. And it's this lack of network clarity—whether along routes ranging up and down the mid and North Atlantic seaboard or across the Atlantic—which is further setting the stage for costly network downtime in the present and near future.

Put another way, today we are facing the increased likelihood that a consumer can simultaneously access Gmail, check Instagram, look at a recipe on Facebook and watch a YouTube video, but be unable to move a $100 out of his or her local bank.

For the large investment bank or brokerage firm doing business intercontinentally, the stakes can be much higher. Consider a scenario where such a global institution relies on two carriers deploying MPLS networks across the Atlantic, lending the organization a false sense of security that such diversity is sufficient.

Now let's envision that a network node fails at a data center located in Ireland, so this same company is forced to reroute traffic out of the UK. Simultaneously, a misplaced anchor from a fishing trawler cuts into a subsea cable, knocking it offline, or further inland a backhoe digs up a terrestrial network segment. In an instant, the investment bank or a commercial trading platform is now entirely cut off between London and New York, potentially costing it millions of dollars in lost revenue.

Lessons Learned Above Sea Level

As with many complex systems, the lesson in the hypothetical network illustrated above is that outages are rarely isolated to single points of failure. Aeronautics and aerospace engineers have long recognized that even small local failures within a complex system can cascade rapidly, accumulate and cause global failure in unexpected ways. One strategy to counter these scenarios is to employ redundancies in the design so that backups or contingencies are in place to prevent a failure from progressing to catastrophic levels. While it's impossible to completely eliminate risk, valuable lessons from failures can be used to continuously improve engineering systems and design processes to ensure that the risks are acceptable.

Consider the lessons learned from Hurricane Sandy in 2012, when millions in the Northeast were without power and several data centers in the financial district were literally underwater. Manhattan is a very complicated system with many core network interconnections running through it, and with much of it is legacy infrastructure prone to failure. These flaws were only compounded because much of the East Coast subsea capacity lands beyond New York City but is also routed into the same complex network entanglement, eliminating what are often perceived as redundant solutions.

Today, we know that a colocation campus, optimally located in New Jersey, which bypasses legacy chokepoints and congested New York City routes—and built 64 feet above sea level, making it impervious to tidal surge—offers a critical solution for multinational businesses needing always-available connectivity and access to data. By strategically intersecting a carrier-neutral subsea Cable Landing Station meet-me room with a Tier 3 carrier-neutral data center, businesses, carriers and financial services institutions can diversify their connectivity options to key hubs across North America, Europe, South America and the Caribbean.

Moreover, as many transatlantic subsea cable systems approach end-of-life in the next several years, creating future chokepoints throughout the Northeast as well as possibly isolating legacy network hand-off points, achieving network redundancy, diversity and resiliency will become even more critical. And then there is the specter of rising sea levels due to climate change. While New York City will likely invest heavily in sea walls, tide gates and pumping stations, these measures won't be able to protect every business, data center or carrier hotel. Resiliency through diversity will become non-negotiable.

In today's business environment, when increased globalization requires the seamless flow of information across borders and between continents, network capability and performance have become the key differentiators necessary for financial services companies to succeed in the connected world. The network is no longer merely the piping and plumbing of the organization, it's the conduit that powers future growth potential, penetrates new markets and reaches new customers.