

# Building High Availability Solutions in Central America and the Caribbean



In the modern computing environment, being always-on isn't a premium feature or marketing fluff. It's a baseline expectation. Business platforms must be available 24/7. That's true no matter the size of the business or its list of clients, to the point a single failure can result in huge losses—in terms of the actual cost of downtime and other, less tangible factors, such as customer experience.

Of course, it's also true that every attempt to improve availability comes with its own cost. These costs rise exponentially the closer we get to 100 percent. This is why the technology industry and its sub-verticals have agreed on certain metrics to gauge how much uptime to expect, such as percentage guarantees: 99 percent uptime, 99.9 percent uptime, and so on. When we calculate these figures over the course of a month or year, we begin to see what they mean in terms of real business downtime:

Availability percentage	Annual downtime	Monthly downtime
99%	3,65 days	7,31 hours
99.9%	8,76 hours	43,82 minutes
99.99%	52,56 minutes	4,38 minutes
99.999%	5,26 minutes	26,30 seconds

Now, let's look at this idea in a more practical sense. As noted, "high availability" typically also means higher costs—exponentially higher, at times—and a more complex network architecture. Achieving a proper balance between cost and reliability means looking beyond the nines and understanding the unique context of your business: How long can the platform you want to protect be inactive without having substantial impact on operations, or the trust of your customers?

For instance, an informational website is not the same as a critical business application. You might hypothetically agree that 99.9 percent uptime is a "sufficient" level of operability for your informational site—in other words, roughly nine hours of downtime per year is acceptable, or at least will not significantly impact business operations. Meanwhile, a critical business app going down for nine hours may well be catastrophic for your business. Thinking about these sorts of things in terms beyond basic 9s and percentages is important to figuring out what downtime truly means to your organization.

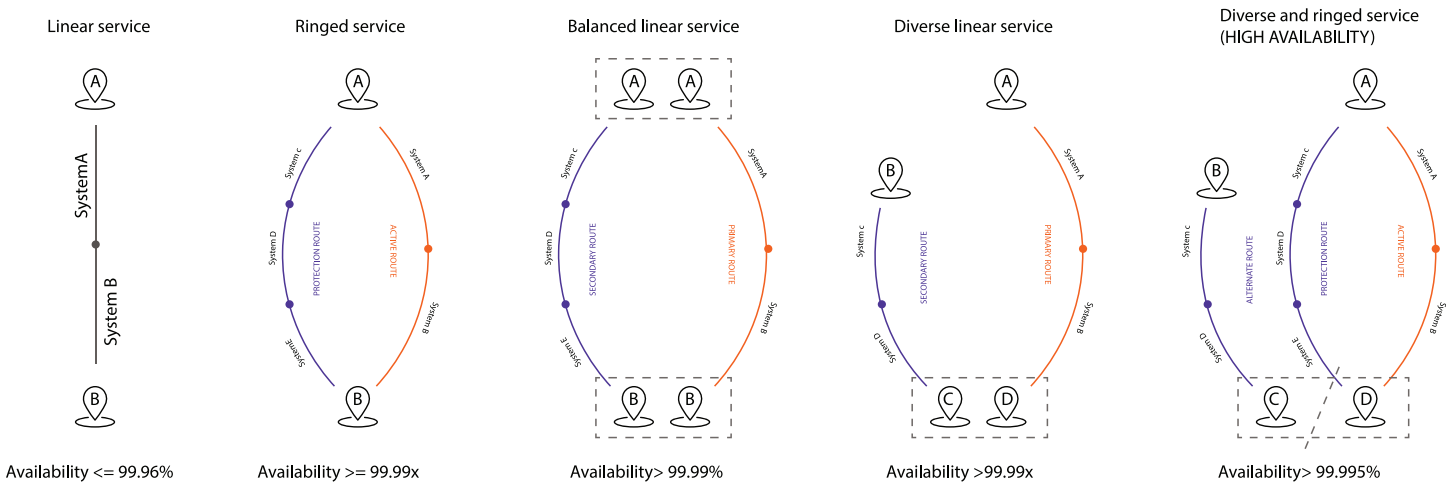


## Fault Tolerance, Networking, and the Path to High Availability

Most of today's business computing systems are distributed. From global startups to established enterprise players, the need to share information and computing resources between sites is more important than ever. That makes connectivity the foundation of high availability for most business.

Because of this, removing Single Points of Failure (SPOF) is critical when working to increase uptime. An SPOF is a node in a system so critical that its failure causes a global service outage—in other words, when it goes down, the whole system does. By contrast, a fault-tolerant strategy looks to identify these SPOFs and make them redundant in the system, thus providing ongoing service continuity even when one link goes down.

At high level, every computer/information system is comprised of three components: compute, storage, and networking. For compute and storage, technologies like virtualizations prevent certain SPOFs by removing the software you need from the hardware it runs on. But more care is needed on the networking side. Let us say a submarine cable connecting an office in Puerto Rico to its servers in Miami is cut by a boat anchor. If that line is the only point of failure, that could mean days without access to critical services. But with a fault-tolerant network, the same organizations can rest easy knowing their critical data is still within reach, with multiple alternate paths available to support it.



These diagrams of various popular network topologies show us the complexity required to achieve an acceptable level of fault tolerance and availability. Every single element creates additional cost and management requirements, each of which must be weighed carefully. It's easy to imagine how connectivity arrangements such as these might significantly add to costs for a company trying to increase availability, especially when considering global enterprises big enough to have presence in several countries.

# How Can Companies in Central America and the Caribbean Achieve High Availability?

Companies in our region have traditionally attempted to build fault tolerance at the network level by hiring multiple different connectivity providers as a means of bet-hedging. But the geographic reality of where we live means this strategy doesn't always produce the desired results.

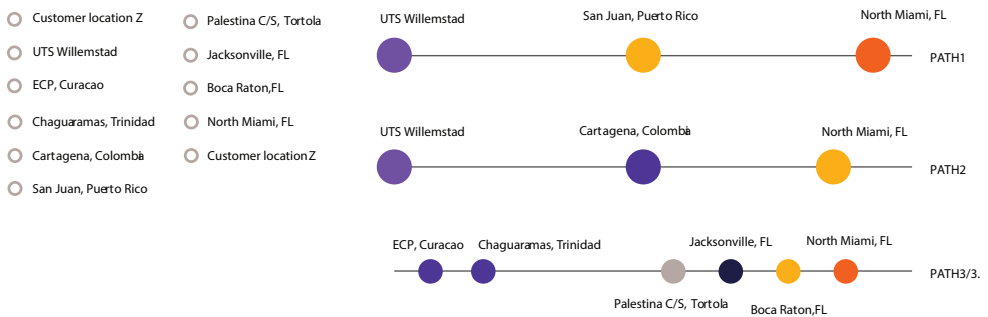
In actuality, there's no guarantee two different providers will use different paths to connect two given points. This is because many of these companies share network segments and submarine cable networks. Moreover, the client typically doesn't have full visibility into the data transmission paths each provider makes available; this makes it difficult if not impossible to eliminate SPOFs with any accuracy. The lesson to take from this all: Multiple providers does not necessarily mean multiple paths for your most important content, data, and workloads.



At Liberty Networks, we're committed to increasing the availability of our clients' critical-mission applications. We own and operate more fiber than anyone else in the region, which allows us to guarantee our customers three independent paths to connect two given sites anywhere in the Pan-Caribbean region, minimizing the impact of a disruptive event and empowering us to offer industry-leading SLAs.

Let's now analyze a typical scenario: a client in Cartagena has datacenters in Miami which must be supported by a high availability scheme. As this diagram shows, our three paths share unavailability points such that one is always in operation, even in the rare event of disruption to the other two!

More, we've got the people to help make the most of our experience and infrastructure. Our people are trained to design and manage networks with high availability in mind, who maintain a principle of diversity of access and transport so we can minimize service disruptions. Because we own the network, we have deep insight into its entire structure, helping us quickly discover incidents and render fixes. On top of that, we offer a singular point of customer service—meaning less finger-pointing and call-hopping if you call with a question or concern.



Whatever your structure or business model, designing for diversity, fault-tolerance, and uptime can be quite challenging. Instead of relying on hazy ownership models and unclear disaster recovery principles for your most important data, it's time to explore a company like us: one that runs the network, knows every inch of it—and understands how to leverage it best for our customers, even in times of disruption.