Couchbase NoEQUAL

DISTRIBUTED CLOUD AND EDGE COMPUTING FOR 100% UPTIME

As cloud computing infrastructures have seen explosive growth over the last decade, and especially over the last year, applications that rely solely on the cloud for data storage and processing are beginning to show signs of strain. This is particularly true of those that require sub-second response times and high availability.

This makes sense because the success of cloud computing hinges on having a reliable connection to a centralized data center via the internet. If the connection goes down, you suffer costly downtime. And even when your applications are up, they must contend with network latency that slows the entire user experience. It takes time for an application to send data to a remote cloud data center for processing and then wait for a result to come back down over the wire in order to respond to an input. When every millisecond can mean the difference between making the right or wrong decision, that's simply not acceptable. Also, in many cases the clients themselves amplify the unreliable nature of internet connectivity by moving from location to location.

A reliable way to eliminate these many risks is to move data processing closer to the applications. Enter the concept of "Edge Computing."

A cutting-edge approach

Edge computing is a term that describes a distributed framework designed to bring data and compute closer to the applications that use them by leveraging servers that are co-located to provide tiers of insulation from network latency or failure. The term "edge" in this case is not a reference to a single end point, rather it describes a layering of data centers in increasingly closer proximity to applications than cloud data centers. This proximity makes applications faster and more available by eliminating dependencies on the internet. There are many other substantial benefits of edge computing that we'll cover below.

When it comes to edge computing, there is a spectrum of models. We'll simplify the types by breaking them into two categories: "standalone" and "distributed."

On one end of the spectrum there's the concept of a standalone smart device where the data is co-located (embedded within the application itself) and always available to the client that needs it. These applications typically run on smartphones, tablets, and rugged devices like the DIAD (Delivery Information Acquisition Device) that all overnight courier drivers carry. The defining characteristic of the standalone model is the notion that the client is the data center, which means the data is captured, processed, and stored directly on the application or device. The data is then usually synced with other data centers, and ultimately to the cloud, as connectivity permits.

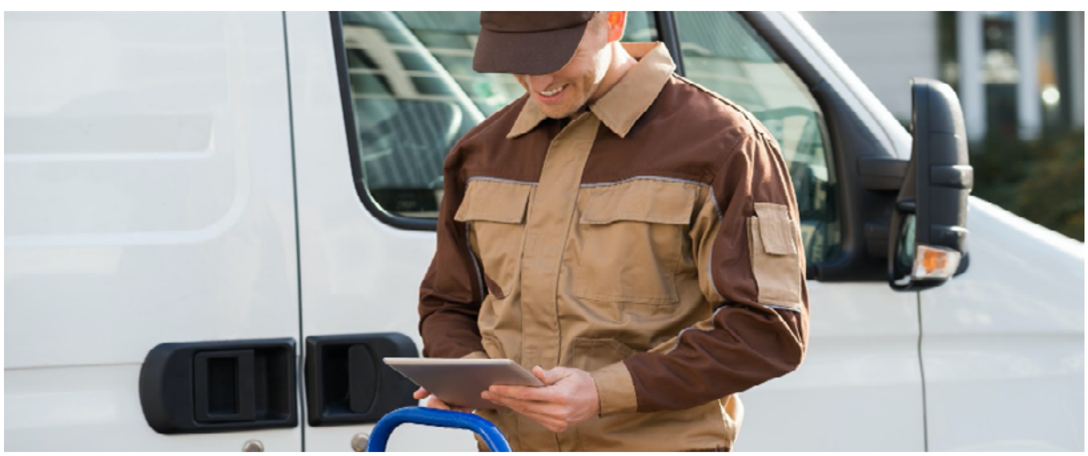

STANDALONE

On the other end of the spectrum you have distributed micro data centers. These "small-footprint" data centers could be in the form of an IT closet in a call center, a server rack in a big box retail store, or on a cruise ship – the footprint really depends on the scale and number of downstream endpoint clients that are being served by the applications. Micro data centers provide a benefit in any situation where high-density, small-footprint compute resources are in demand, such as retail point of sale (POS) systems, IoT applications, content delivery applications, and applications that process data for AI and machine learning. They're especially beneficial in environments that experience no internet connectivity for extended periods of time, such as on a cruise ship or airliner. Micro data centers power the distributed cloud architecture and enable businesses to get the benefits of the cloud while also bringing compute and storage resources on premises to increase responsiveness and reduce downtime.

DISRIBUTED

From the standalone to the distributed ends of the spectrum, the one constant is that data and compute are brought closer to the application or client as a mitigation to latency and downtime. The edge computing model presents unique benefits and overcomes many challenges of cloud-only architectures.

The key benefits of edge computing

As a model for minimizing the risks of cloud computing, the distributed edge computing architecture presents a best-practice approach for tomorrow’s technical innovations by providing these four key benefits.

Reliability and resiliency

Edge computing provides guarantees of business uptime through uninterrupted access to data regardless of internet connectivity. Having a reliable method of data access that can withstand inevitable connection issues helps better ensure application performance and, in the case of consumer applications, a better customer experience, which is critical for loyalty and repeat business.

Speed

Cloud computing offers benefits of scale, easy management, and cost savings, but in many cases the data is simply too far removed from the application to facilitate real-time performance. Trying to achieve real-time speed solely by optimizing the network is simply not feasible from a cost or even a technology standpoint. Applications such as those used in healthcare or manufacturing rely on a high volume of in-transit data. Low latency, sub-millisecond access times are critical, and regardless of internet bandwidth, such applications simply can’t be served with a cloud-only model. Moving computing to the edge is the only way to gain the necessary speed these mission-critical applications require.

Security
Data privacy, governance, and regulatory policy compliance are imperatives for applications that handle sensitive data. Healthcare applications, for instance, require that sensitive information must never leave the premises. And consumer applications such as fitness trackers make recommendations based on lifestyle and location data, but that information must be processed locally because consumers don't want it stored or processed in the cloud. A key value point for edge computing is that sensitive data never has to leave the edge – it can be stored within the appropriate micro data center tier only if and when it is required.

Bandwidth usage
Cloud computing delivers savings from an efficiency standpoint. Moving more computing to the edge creates additional savings by reducing bandwidth usage and by reducing load on cloud servers by lowering the volume of data transferred over the internet. Costs also become more predictable because you can be more precise in estimating bandwidth requirements based on the distribution of edge devices and the expected throughput.

Let's take a closer look at why the distributed cloud model is being hailed as the architectural approach for tomorrow's tech breakthroughs.

Cloud formation

Over the last decade, enterprises have been shifting their technology infrastructures from on premises to the cloud because the cloud brings economies of scale, the efficiency of infrastructure manageability, and the elasticity to adjust compute power in proportion to demand. All these benefits also equate to cost savings and a more nimble business model.

Cloud offers an economical option for enterprises, but it also relies on internet connectivity, which introduces a point of failure over which there is little control. And the inherent latency in sending data to and from the cloud becomes a hindrance to real-time requirements when sub-second response times are critical. A business that loses the connection to its critical data can lose millions of dollars in an instant. In 2018, the retail giant Target suffered a two-hour outage of its POS system nationwide, which cost the company an estimated $50M in lost revenue – again, this was in just two hours! The economic impact was further exacerbated by negative consumer opinion on social media that eroded Target's reputation for months after the incident. In short, relying solely on the cloud computing model to power business-critical applications is a surefire way to suffer costly business downtime.

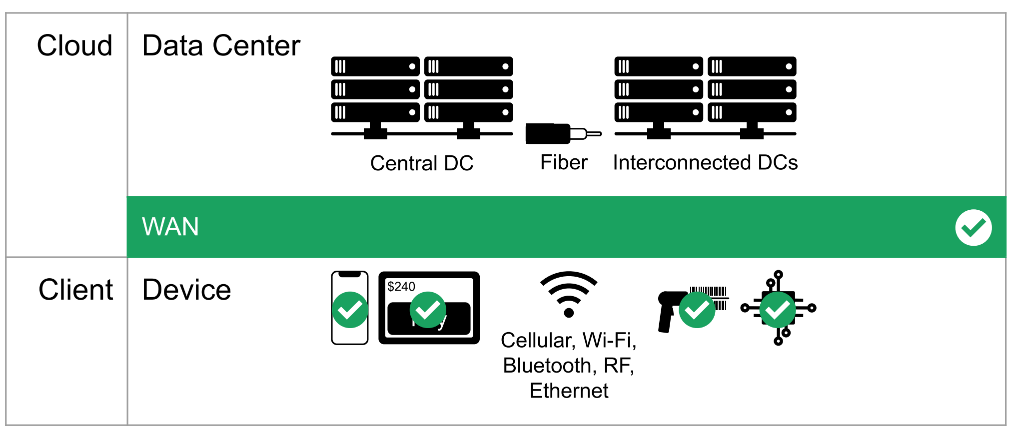

Figure 1. When the cloud data center is available on the WAN (wide area network), client applications run as intended.

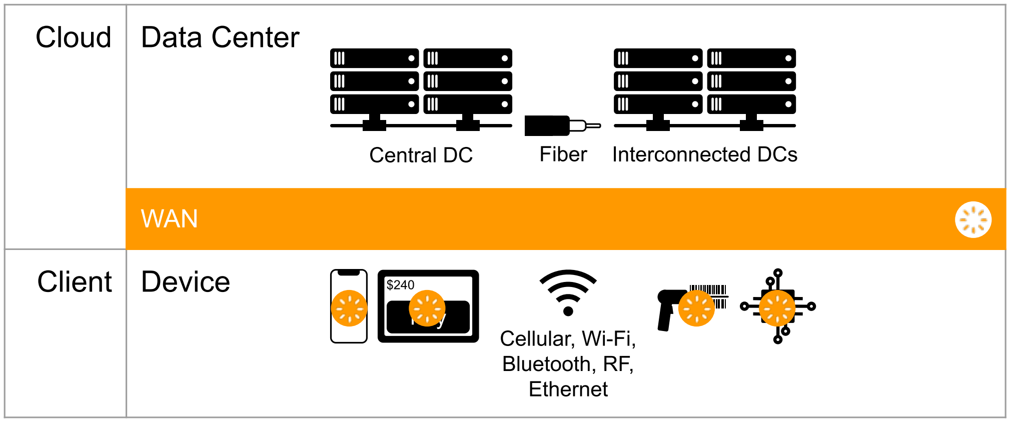

Figure 2. When internet connectivity suffers latency, client applications slow down.

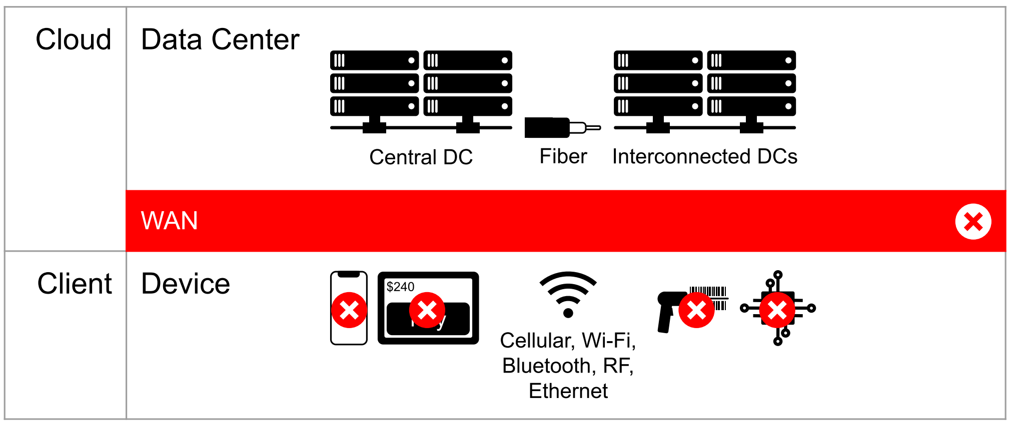

Figure 3. When internet connectivity is lost, client applications cannot access or process data and grind to a halt, resulting in business downtime.

As figures 1-3 illustrate, cloud computing's reliance on internet connectivity presents a potentially catastrophic point of failure that can severely hamper applications that require real-time responsiveness and 100% uptime. In such cases, a distributed edge computing model is the better option.

Distribute the processing, eliminate the downtime

For applications that require faster responsiveness and better business uptime, the challenges of cloud computing can be overcome by capturing and processing data at the edge.

In the following figures you'll notice that edge computing still involves syncing data via the internet (WAN) to the cloud for eventual data storage and aggregation across the entire environment. But edge computing layers data acquisition and processing closer to devices and applications via micro data centers available on a LAN (local area network) or PAN (personal area network) to speed up processing and eliminate points of failure.

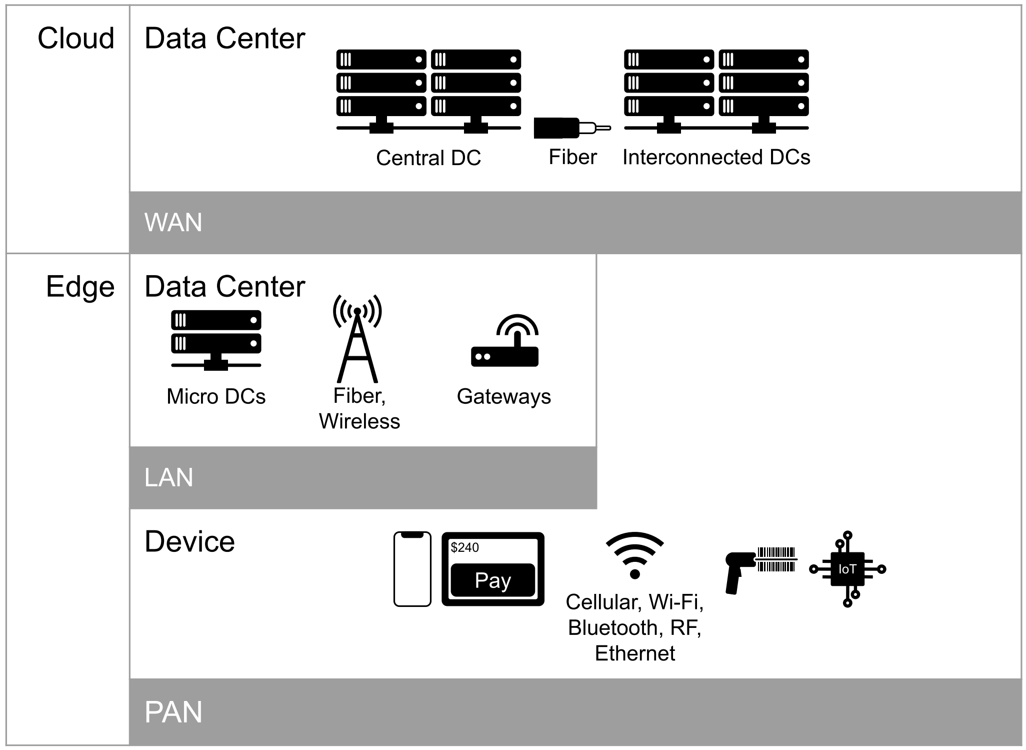

Figure 4. The distributed edge computing architecture.

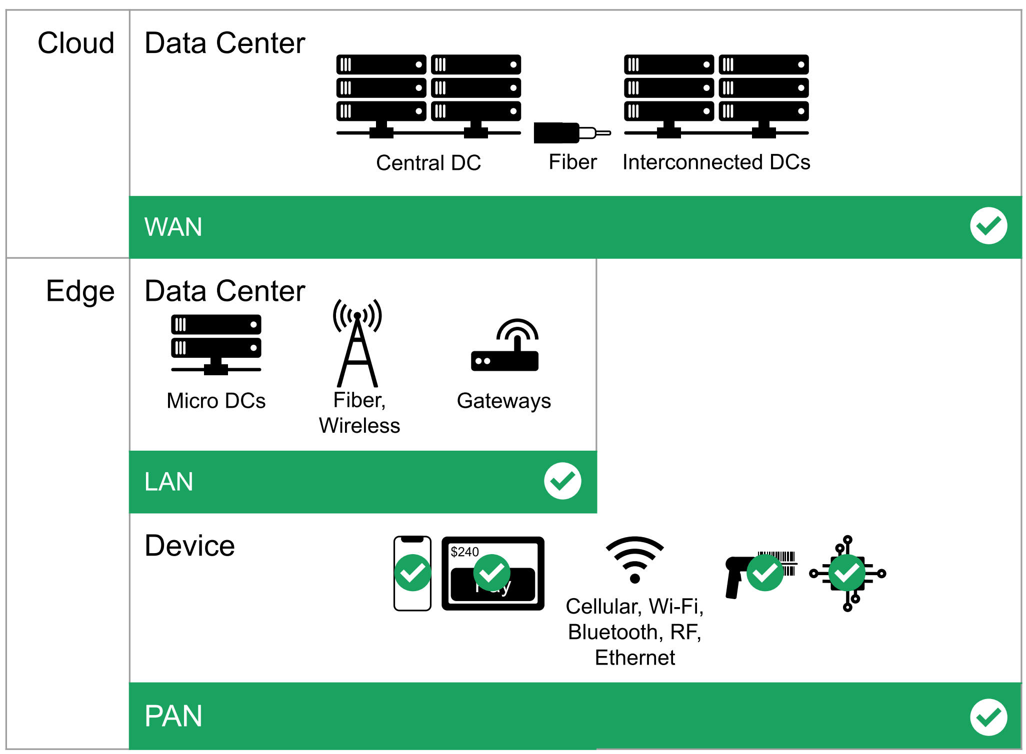

Figure 5. When the WAN/LAN/PAN are available, applications run as intended.

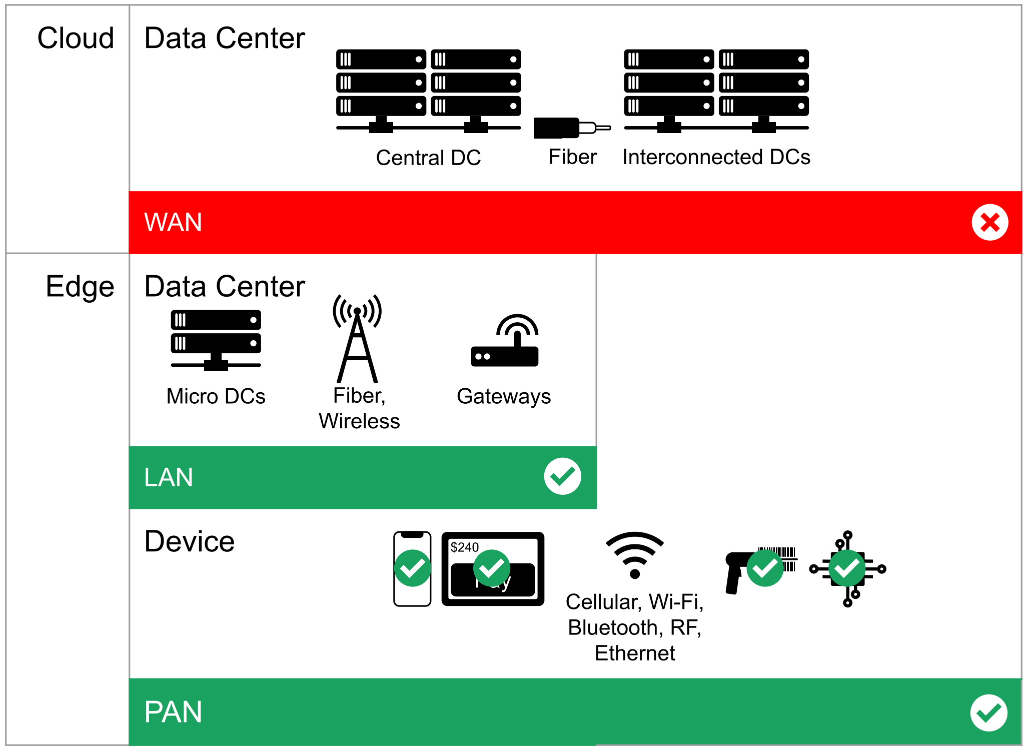

Figure 6. When the WAN becomes unavailable, applications continue to run as intended by processing data in micro data centers at the edge in the LAN/PAN layers. When the WAN becomes available again, applications sync to the data centers in the cloud.

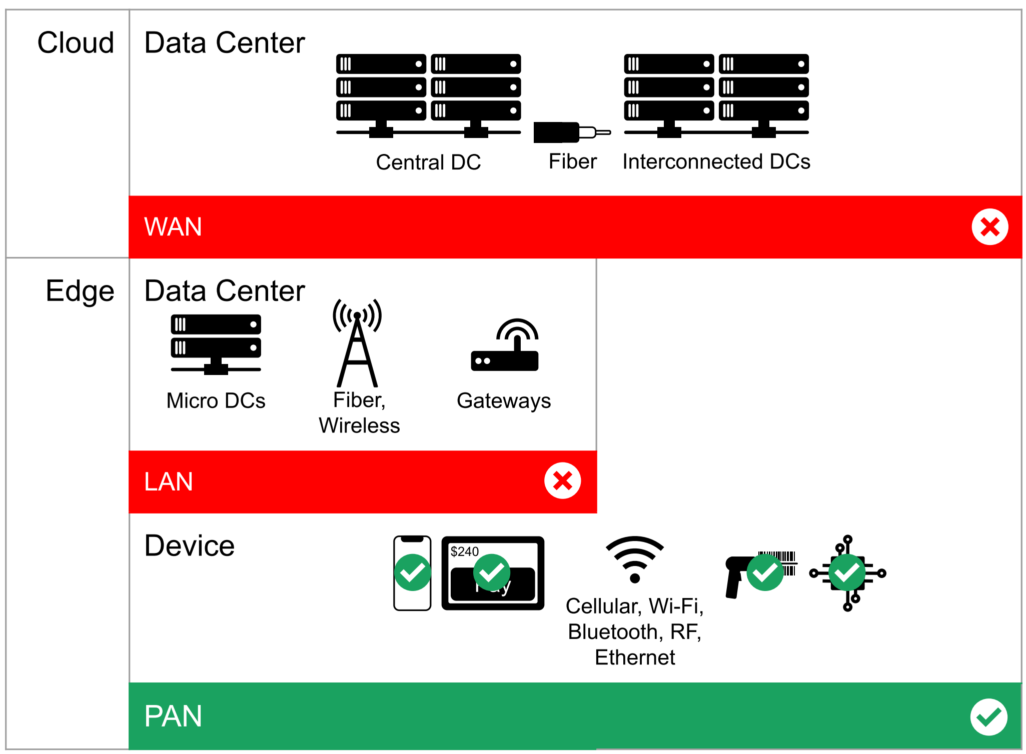

Figure 7. If the cloud data center and LAN micro data center become unavailable, applications continue to run as intended by processing and syncing data directly on devices that act as micro data centers.

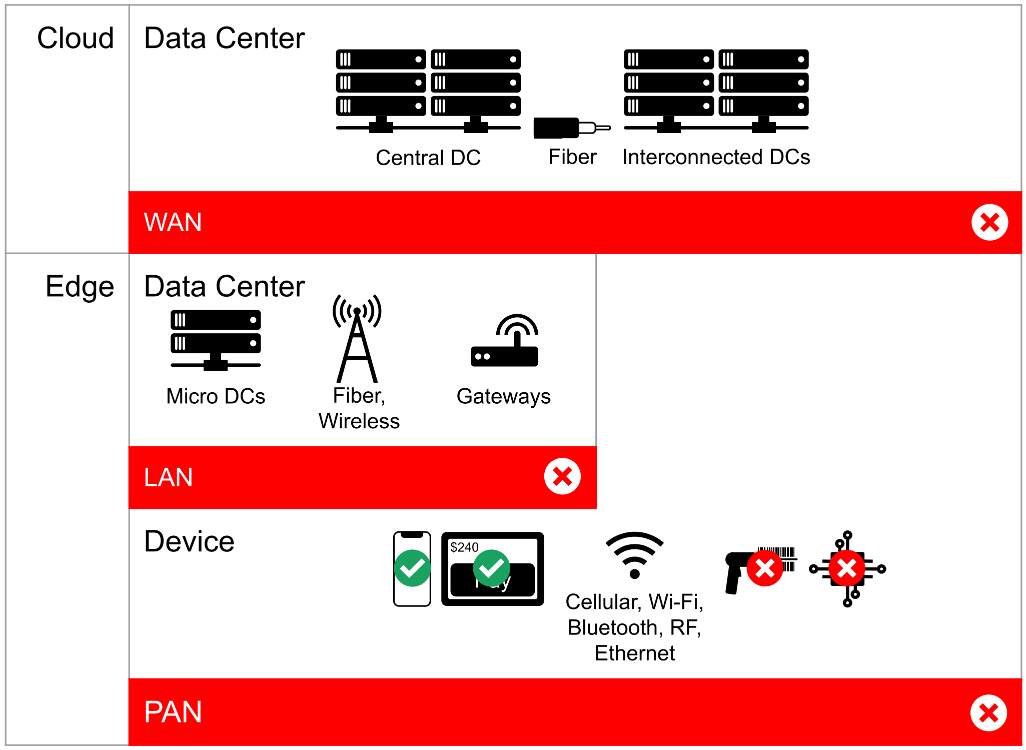

Figure 8. If all network layers become unavailable, devices and applications with embedded lightweight data storage can serve as their own micro data centers and continue to run as intended until connectivity is restored.

As figures 4-8 illustrate, by distributing data processing to data centers located closer to the applications and devices that consume it, the edge computing architecture mitigates risks of latency and business downtime.

A framework for today's critical applications and tomorrow's innovations

The distributed edge computing architecture presents a model for future innovation through a simple concept: bring data closer to the clients that use it, which enables a new class of applications that react to inputs instantly and never go down. Edge computing has the potential to revolutionize industries such as retail, hospitality, and finance by powering a premium "always-on" real-time consumer experience. It can dramatically impact healthcare and emergency response through its ability to instantly and reliably help with life-or-death decisions. And it can greatly increase the effectiveness and safety of teams operating in isolated, low- or no-internet environments such as remote wilderness, sea-going vessels, mines, or even in space, where losing access to critical data is simply not an option.

In order to take advantage of the distributed model, developers should strive to leverage technology designed to facilitate edge computing, such as lightweight data storage built to be embedded directly into applications and devices, as well as built-in synchronization capabilities to ensure that data is always available and never lost or corrupted. Using edge-native technologies makes development faster, easier, and less costly, and it allows you to focus on the core competency of your application without worrying about speed and connectivity issues.

A distributed cloud architecture is rapidly becoming more than a competitive advantage, it's an imperative best-practice approach to modern computing that enables a new class of resilient, fast, secure, and efficient application for 2021 and beyond.

About Couchbase

Unlike other NoSQL databases, Couchbase provides an enterprise-class, multicloud to edge database that offers the robust capabilities required for business-critical applications on a highly scalable and available platform. As a distributed cloud-native database, Couchbase runs in modern dynamic environments and on any cloud, either customer-managed or fully managed as-a-service. Couchbase is built on open standards, combining the best of NoSQL with the power and familiarity of SQL, to simplify the transition from mainframe and relational databases.

Couchbase has become pervasive in our everyday lives; our customers include industry leaders Amadeus, American Express, Carrefour, Cisco, Comcast/Sky, Disney, eBay, LinkedIn, Marriott, Tesco, Tommy Hilfiger, United, Verizon, as well as hundreds of other household names. For more information, visit www.couchbase.com.