



Life at the Edge: An Edge Computing Primer

Edge computing is quickly becoming important to organizations managing voluminous and constant data feeds with stringent performance requirements. As more companies consider the role of the edge as a solution to the challenge of moving this data all the way to the cloud and back, there has been confusion about edge services and how to leverage them.

Below are some of the most common questions enterprises have about the edge, including which use cases are best suited for edge computing and how it relates to broader cloud goals.

What is edge computing?

Before we get into more details around the use cases, we have to know what, exactly, the edge is. [According to IDC](#), the edge consists of a “mesh network of micro-data centers that process or store critical data locally and push all received data to a central data center or cloud storage repository.”

In plain English, edge computing consists of smaller data centers in locations closer to where users are, letting them communicate with data without latency challenges, to deliver optimal performance and efficiency for all data.

Why are edge services of interest to enterprises?

Edge-based solutions help with a couple of serious challenges:

Latency

Any time you try to move data across long distances – like to the cloud – it takes time. While this is not an issue for some applications, it is for others. Workloads and datasets that are highly transactional or interactive don't function as well when they are far away from users.

Upload versus download

This relates to the very design of the Internet. Users and applications generate enormous amounts of data, a trend that is only continuing. Getting all this data uploaded to the cloud presents a dilemma because the Internet itself is optimized for downloads. Trying to upload this data is akin to having one highway lane for uploading versus five lanes for downloading.

The edge, then, aims to solve these dual problems by extending the cloud – especially for data access – closer to users, where it can be consumed and managed, as if it were stored locally. This lets applications take advantage of all the cloud has to offer, but without any of the drawbacks.

Why has it been hard for enterprises to migrate to the cloud?

In the past, organizations have expressed security concerns about storing their most sensitive data, especially customer

data and intellectual property, anywhere but in their own data centers. There have also been concerns about compliance and regulations. These issues have largely been solved by cloud providers, especially with the introduction of features like built-in backup and disaster recovery.

The latency and upload problems noted above explain some of the reluctance enterprises have about going to the cloud. This presents a simple fact: there are many applications and corresponding datasets that will never be an exact match for the functionality the cloud provides.

This doesn't mean that on-premises is the best option, though. When you examine a lot of the apps that aren't an exact fit for the cloud, you often come to the conclusion that they aren't the best fit for on-premises, either. Costs can become prohibitive and organizations have come to believe there is a better way.

This Catch-22 illustrates why the ideal strategy isn't the edge or the cloud; it's the edge with the cloud.

The best use cases for edge computing

The use cases for edge computing are growing by the day. Most have one (or more) of these things in common:

No latency allowed

We're not all neighbors with the public cloud. Public cloud data centers are often located in remote locations to get the benefits of inexpensive land, which keeps costs down. The problem this creates is added latency, due to the speed of light limitations, switches, routers and servers.

Sending data to a cloud data center for analysis and then bringing it back to act on locally isn't an option for certain applications or datasets. Operational applications tracking and responding to critical events in real time, financial applications that process millions of dollars in seconds or healthcare applications where patient data is needed instantly are good examples, but there are many more.

The edge eliminates latency by keeping data close to users, making it a great fit for these apps.

Machines talking to machines

With the advent and growth of the Internet of Things (IoT), there are many applications that never interact with people before executing a command. A common example is something like the connected, and eventually, self-driving car. These vehicles have to communicate with each other and with the operational hubs at an extremely fast rate, using a large amount of data. Clearly, shipping that data to a far-off cloud for analysis simply doesn't work – milliseconds can mean the difference between an accident and safe travel. Likewise, the need for virtually 100 percent uptime and serious computing power makes it a poor fit and cost-prohibitive option for on-premises.

You don't need to be talking about something as big as fully

connected and autonomous cars, though. Any application that requires rapid, seamless communication between machines or devices, with little or no human input, is a great case for the edge.

An excess of data

What the previous two use cases largely have in common is they generate a lot of data. Due to the download-optimized and unpredictably latent Internet, virtually any app that generates a lot of data and has to act on it quickly, is not a fit for the cloud. As more people and apps upload data (photos, videos, etc.), the upload stream quickly becomes a bottleneck.

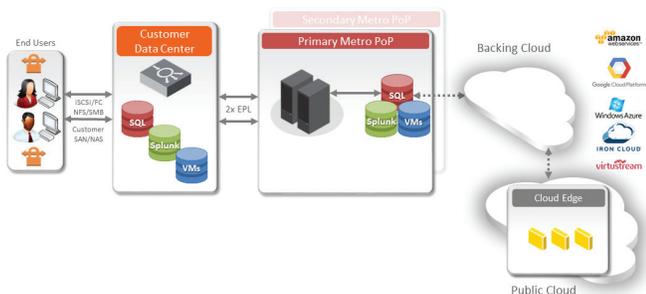
The edge, however, doesn't rely on the Internet. It processes uploads and downloads equally fast, eliminating potential upload delays.

The edge in partnership with the cloud
Any data that needs to be analyzed and acted on immediately is great for the edge, but what happens once the data isn't actively used anymore?

Storing older data at the edge isn't necessary or cost effective. On the other hand, a lot of it needs to be kept for record keeping, compliance or any of a number of reasons. Why not move that data from the edge to the cloud for long-term storage and on-demand access? It can still be accessed easily and without egress, but you can now realize the cost benefits the cloud has been trying to offer for some time.

This switch can even happen automatically, based on user preferences and usage patterns. "Hot" and "warm" data is stored at the edge – available with flash performance for on-demand access – in a point of presence close to your location. "Cold" data that just needs to be kept for posterity or compliance can be stored in the public cloud and be accessed readily, without egress fees.

The ClearSky Architecture



The ClearSky service

ClearSky is serious about the edge. Our fully managed service that includes on-demand primary storage, offsite backup and DR, makes it easy to get your apps and data to the edge, and optimize them once you're there.

VMware integration

One of the biggest stumbling blocks for organizations going to the cloud is migrating their apps and data. Our complete integration with VMware makes this process easy.

On-demand access

As your needs change, your storage footprint can change right along with them -- no provisioning, installing or managing equipment. ClearSky's managed service lets you expand or contract your capacity quickly, easily and securely.

Edge-based network

Our edge-based network makes use of the edge without your having to think about it. With our fully managed service, your data is optimized across its entire lifecycle using our patented Smart-Tiered Caching™, keeping "hot" and "warm" data close to users and all data backed to the public cloud of your choice.

"Hot" and "warm" data are cached at the edge, in metro PoPs, for all those applications that need data quickly and seamlessly. All your data, including "cold" and archived data is stored in multiple cloud locations and easily accessible whenever you need it.

Backup and DR built-in

ClearSky also provides built-in, offsite data protection, automatically. This means no replication, no backup licenses, no DR licenses, no maintaining secondary data centers, and no need to provision resources to manage it all. You simply get backup and disaster recovery as part of your storage solution at no added cost.

This model provides a number of important benefits. First, with hot and warm data kept at the edge, if an outage occurs recovery is virtually instant -- allowing for recovery time objective (RTO) and recovery point objective (RPO) of 0.

Reduce costs by 50%

This as-a-service, edge-based approach allows customers to gain the benefits of the cloud while minimizing on-prem footprints and investments and reduces costs by more than 50 percent in the first year alone.

Additional edge resources:

- [ClearSky Data Edge and Cloud Solutions](#)
- [Edge Computing Roundtable: Friend or Foe of the Cloud?](#)
- [Edge Computing: Is It for Everyone?](#)
- [The Edge or the Cloud? It Depends on the App](#)
- [4 Factors Driving Companies to the Edge](#)

Have more questions, or want to talk? [Contact us.](#)