Getting Closer to the Mission: How Edge Computing is Transforming Federal IT

By: Carl Conner, Director of Enterprise Solutions

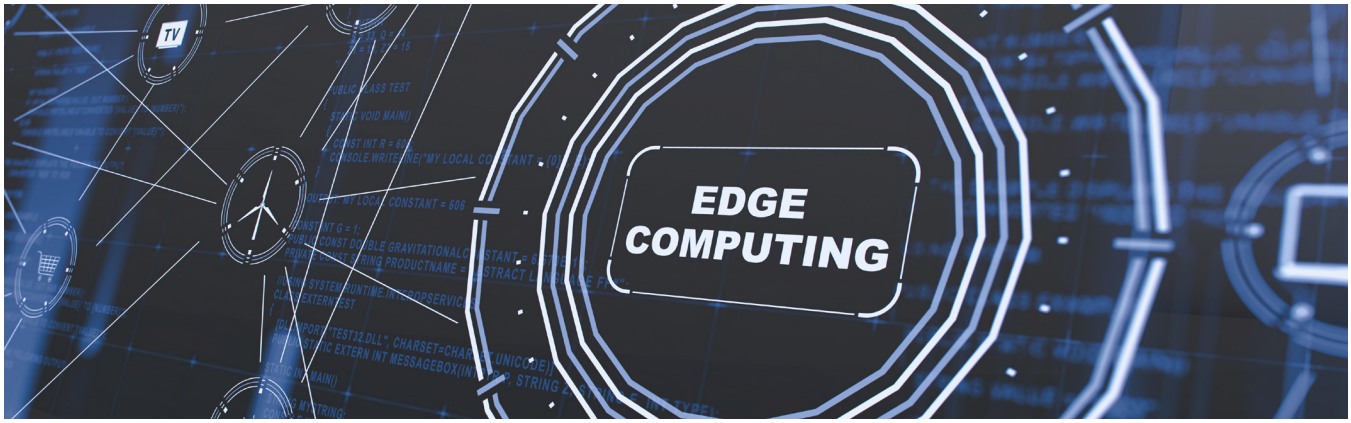

The explosive growth of data and access to real-time applications such as video processing and artificial intelligence is transforming the way the Department of Defense and federal civilian agencies achieve their missions. Cloud computing paved the way for more rapid scalability to support this new landscape; however, the very nature of remotely located cloud servers implies a long round trip for data. In an environment where "real-time" data processing is no longer real enough, edge computing is quickly becoming an invaluable technical resource across the federal government.

What is edge computing? Simply put, edge computing is data processing that is located close to the edge – or where things and people produce or consume that information (Gartner). In this article, I'll review the four major types of edge computing, their key benefits, and how they can be applied in real-world settings in the federal government.

Note: As with any technology, things are changing rapidly. Some of the terminology used in this article may change over time.

Mobile Edge Computing (MEC)

The first major type of edge computing is Mobile Edge Computing (MEC). MEC is a highly distributed computing environment where applications are deployed, and content is stored and processed, near mobile users. You can think of this as a "downward push." Hsu, Wang, Zhang, and Kobusinska (2018) describe MEC has having three key distinguishing features: close physical proximity to users, support for mobile devices, and a dense disbursement of MEC servers in a given geographical area. MEC resources are often based in LTE base stations or collocated in and with wireless access points.

For example, Twitch, a popular streaming site for video games, might run a promotion in Las Vegas for a new game. In preparation, local servers are optimized so that large numbers of users can download and stream the game on their mobile devices without interruption. Another, more relevant example in the federal space, is mobile law enforcement. Imagine U.S. Customs and Border Protection is running an operation in the destitute desert of New Mexico. Using MEC, static GIS mapping data could be cached locally where fast changing data such as the location of possible targets could be overlaid live at the mobile level, providing for faster and improved access.

MEC has several key benefits:

1. Lower latency for data pulls from the end user
2. Increased security
3. Better overall user experience (UX)

MEC servers have also been successfully used as extended Content Delivery Networks (CDN). For example, users in a specific city might be more interested in video streams for local sports teams or local traffic conditions, but not the real-time traffic data for a city hundreds of miles away. MEC provides a way to provide this content locally, which reduces network traffic back to the cloud servers, increases security, and can even save energy. Additionally, MEC can be used on certain devices to support opportunistic offloading of computational resources for CPU intensive tasks, such as near real-time voice recognition (i.e. medical transcription) and translation (i.e. Google Translate). If the compute power of a mobile device is limited, it may be worthwhile to offload the computation of these functions to a MEC server and then retrieve the result. This makes particular sense for areas of known high density use, like a busy international airport.

Fog Nodes/Fog Computing

Fog computing versus cloud computing—it's an analogy in a name. Clouds are remote and seldom seen or touched. Fog, on the other hand, is dispersed among us. Therefore, fog computing is compute and storage resources that are located geographically close to the end user or collection device.

Fog Computing was first proposed by Cisco in 2012 as a viable third tier that would rest between mobile devices and true cloud servers. In their white paper titled Fog Computing and the Internet of Things: Extend the cloud to where the things are Cisco explains that "analyzing data close to the device that collected the data can make the difference between averting disaster and a cascading system failure" (Cisco, 2015). This is especially useful when you have a "data tsunami."

Take for example the Large Hadron Collider (LHC) at the European Organization for Nuclear Research, known as CERN. The LHC produces over 500,000 Petabytes of data per day (yes, Petabytes). But most of this data is not needed and will quickly overwhelm storage systems. Fog nodes are used to quickly throw away 99.9999% data that is not relevant to the experiment in question, and only send on for storage those bits that are worth further examination.

Fog computing has several key benefits, including:

- Superior data filtering and verification
- Low latency
- Reduced network traffic back to the primary cloud data center

Micro Data Centers (MDC)

The third type of edge computing is Micro Data Centers (MDC), otherwise known as a "private cloud in a box." MDCs are typically set up in an established IT environment, taking up only a single or half rack depending on compute and storage needs.

MDCs are ideal for smaller, remote environments such as aboard aircraft, ships, or in mobile trailers. They provide the compute and storage benefits of the cloud locally, while remaining separated from the larger cloud infrastructure back at "headquarters." This separation can be temporary or permanent, depending on mission needs.

MDCs come in sizes from 1 to 100 kW and thrive in environments with an intermittent connection to the Internet. They must work in isolation such as in-flight or through slower Internet connections such as a Trusted Internet Connection (TIC). For example, a Boeing 787 with a full complement of IoT sensors generates over 0.5 TB of data per flight which is captured by an MDC in flight and then uploaded to the cloud once the aircraft is back on the ground. Other use cases include digital imaging in military forward hospital units, mobile command centers used during emergency response, or within classified environments such as SCIFs.

Cloudlets

The final type of edge computing is a Cloudlet. A cloudlet is a hardened MDC that is designed for deployment outside of a data center or traditional IT environment and is typically positioned near a wireless tower or wireless access point (WAP).

It is important to note that "hardened" can mean different things to different customers. For some it might mean being hardened against weather elements. For others, like military personnel on the front lines, it might mean against explosive devices.

A good use case for Cloudlets is computational offloading for mobile devices. Allowing a Cloudlet to quickly process something like video evaluation and return a result to a mobile device can not only improve user experience, but also dramatically improve battery life. Ma, Lin, Zhang, and Liu (2018) did research on Internet of Things (IoT) sensors offloading computational requirements to nearby Cloudlets. Maintaining Quality of Service (QoS) is challenging for many IoT sensors due to overloading of networks. Cloudlets can assist with data filtering and aggregation. Take for example, a security agency that is monitoring 20 physical locations around the country each of which has 250 sensors reporting their status once per second. That would create 432 million data points per day that need to be backed up to the cloud. If a Cloudlet is used at each location, it can monitor, collect, and aggregate that data by sending one single summary report every five minutes for normal operations and then all data for the site if it detects a break in. The majority of the time the data is of limited use after a short time period and can be safely discarded from the cloudlet after an agreed upon time frame.

Lee and Lee (2018) provide examples of using Cloudlets to fix latency issues with a mixture of small cloud servers and MECs near the end user. They explain that key to being able to solve latency issues is using a hierarchical Cloudlet model. In their models, MEC severs are used as a first abstraction layer closest to the users. These MEC servers are backed by and aggregated into Cloudlets which then provide the final connections where needed into cloud data centers (this assumes that there is a logical hierarchy in place that can be built upon in advance to anticipate needs).

Key Challenges and Limitations of Edge Computing

Edge Computing has incredible potential for the future but is still limited by some substantial constraints.

First, Edge Computing will always be a collection of niche products. When building a network, Edge Computing will generally be something added on as a third layer to a design in specific areas that the two layers of end point and cloud need some assistance. Edge Computing is limited in effectiveness if it is deployed without a backing cloud layer behind it (although this can be done as a private cloud solution).

Additionally, Edge Computing solutions simply cannot scale in the same way that true cloud systems can. The concept of rapid elasticity for cloud is limited to the amount of pre-planned hardware at a given geographical location. Adding additional virtual machines is possible up to a pre-planned level, but expansion beyond that point is more painful than adding more VM's in a datacenter or even bringing another data center online. This is most critical in areas with high population density or large fluctuations in density (Ma, Lin, Zhang, and Liu, 2018).

Finally, heterogeneous networks are a challenge for offloading tasks in Edge Computing (Park and Lee, 2018). This is due to several reasons, but first and foremost is the mixture of competing interests in privately held data usage. Simply put, everyone wants to be first and there are few agreed upon QoS standards that tell all cell phones and tablets to let a packet from a first responder through in front of a request for Yelp restaurant reviews.

The U.S. Military: An Opportunity for Edge Computing

Military environments are an excellent fit for Edge Computing. The U.S. military has great control over the hierarchy of their IT environments. In fact, although not specifically mentioned by Lee and Lee (2018), the Cloudlet hierarchy approach fits military data model quite well. Deployed units in foreign arenas have a command and control structure that rolls up neatly into increasing levels of complexity that could match nicely with Fog and MEC Computing that is aggregated into Cloudlets and then back hauled to cloud servers near mission commands in the Continental United States (CONUS).

The military is also in a unique position to solve or at least avoid the challenges of heterogeneous networks. They can purchase homogeneous end devices for several classes of devices such as smart phones, tablets, and IoT sensors. Once purchased, they can enforce software provisioning that ensures a higher level of security than is practical in a civilian, Bring Your Own Device (BYOD) ecosystem. Antal (2016) explains that there is an increase in the use of homomorphic encryption (HE) for military use.

This is often an impractical level of encryption by small devices but could be performed at the Edge Computing level before transmission back to a CONUS data center.

Additionally, the military can, in fact, programmatically enforce QoS rules on its own devices for different types of traffic that might otherwise overwhelm a Fog Node or Cloudlet.

Another reason Edge Computing is ideal for the military, is that the military has the budget that is required for purchases at the scale needed to make Edge Computing practical. The Department of Defense has been dramatically increasing the expenditure on cloud and cloud-related services in recent years. For example, the DoD issued and awarded an approximately $10 billion Joint Enterprise Defense Infrastructure (JEDI) contract for a single, unified DoD Cloud solution. This will be the one of the single biggest purchases of IT support for the US military ever.

Finally, there is a movement to move the hypothetical into the military cloud. Antal (2016) provides an interesting insight into the NATO military approach to early Cloud adoption. He describes an ongoing effort to provide Modeling and Simulation as a Service (MSaaS) for NATO. This is considered to be a rapid way to test out scenarios and run what if explorations.

Conclusion

It's no secret that edge computing is quickly becoming a "must have" in the future of technology. It offers a middle tier between mobile devices and the cloud, delivering lower latency, increased security, and better enablement for mission success.

At Compass Point Federal, we are committed to helping government agencies modernize their IT infrastructure and systems to drive innovation and value. Our team of experts deliver highly specialized, data-centric services across government enterprises.

Compass Point is a proud 8(a) contract holder on the ITES-3S contract (W52P1J-18-D-A105), an Indefinite Delivery Indefinite Quantity (IDIQ) contract that includes a full range of services and solutions necessary for the Army to satisfy its enterprise IT requirements.

If you're interested in learning more about how edge computing can benefit your agency or Department, visit:

WWW.COMPASSPOINTFEDERAL.COM