Empowering video surveillance and video analytics with a distributed cloud

A Futurithmic blog by Monica Paolini, Senza Fili

July 2020

Distributed cloud (and specifically edge cloud), has a measurable advantage over a monolithic centralized cloud -- especially when it comes to use-cases involving high-volume traffic, tight latency requirements, heavy processing requirements, and location-specific data, like video surveillance.

Video surveillance and video analytics powerfully illustrate the need, as well as the challenges, in pushing some functionality towards the edge of the network.

Video is a mature, but challenging service: meet the requirements, capture the opportunity

Not surprisingly, video surveillance and analytics is the enterprise 5G use case that CSPs find most compelling. In a global survey by Nokia and Omdia, 66% of CSPs list it as one of the top four use cases for 5G. The appeal comes from applications that are mature and widely deployed across virtually all verticals, but with demanding requirements that solutions like wireline (e.g., Ethernet) or wireless (e.g., Wi-Fi and 4G) technologies cannot always meet.

Wireline connectivity may not be available or may be too expensive in some locations; or cameras must be mounted on mobile assets and vehicles or carried around by employees. Legacy wireless technologies may not meet the latency, jitter, throughput, and reliability requirements that some video applications require.

Because cameras can easily be connected to 5G networks, video surveillance services on 5G can be deployed right away. The value proposition is well proven. Deutsche Telekom and Nokia held a 5G trial at the Port of Hamburg, the largest port in Germany. With 1,800 employees handling 9 million containers per year and over 10,000 trucks a day, it shows how video surveillance and video analytics can work together to track assets, monitor safety, and manage operations using a low-latency wireless network with good coverage and mobility support.

Shifting the heavy work to the edge

5G brings lower latency, higher speed and capacity, and better traffic management. But the edge cloud is going to be a necessary complement to make 5G shine in many environments. In the 5G network in the Port of Hamburg, there was a 50% reduction in round-trip latency, as processing moved to the edge (3 km from the port) from a more central location in Nuremberg, 500 km away.

This is what the edge infrastructure brings to the table:

- **Lower latency.** 5G lowers the RAN latency, but a major component of the end-to-end latency – and hence the latency that the user or the application sees – is what happens beyond the RAN, as the data has to move across the network. The closer to the RAN the cloud infrastructure is, the lower the latency is. For many Ultra Reliable Low Latency Communications (URLLC) video analytics applications, the edge infrastructure is going to be a non-negotiable requirement.
- **Location-based processing.** Most video surveillance and analytics data is collected and used locally. If a camera spots suspicious activity, security staff on the ground need to be

alerted, not management thousands of miles away. As long as applications can be remotely managed, keeping processing local increases their efficacy.

- **Lower transport needs and costs.** Processing and storing video at the edge dramatically reduces the amount of data to be transported across the network. For instance, only alerts and reports may need to be sent to the enterprise headquarters.
- **Higher security and reliability.** By keeping traffic closer to the source, the edge cloud gives enterprises and CSPs a higher level of control over how to manage specific security and reliability requirements.

It is more than the edge: Negotiating tradeoffs

This doesn't mean CSPs move all processing and storage to the edge. In addition to edge cloud, we also need a distributed cloud that makes it possible to assess the tradeoffs between centralized, regional and edge infrastructure.

The edge cloud almost invariably reduces latency, but some applications do not require URLLC or extremely low latency, so the regional or centralized cloud may meet the requirements. 5G supports extremely low latencies, but that typically comes at a cost. CSPs must carefully assess the cost/performance tradeoffs.

Space, location, and power availability may be not available or too expensive at the edge, especially at some on-prem enterprise locations. Some CSPs have good availability of real estate (e.g., central offices) that can be cost-effectively used at the edge or regional locations, and they may want to use those locations even for applications that do not have the tightest latency requirements.

The more distributed the cloud is, the more complex it becomes, simply because there are more cloud elements to manage. Automation, real-time optimization, and orchestration take a more central role to ensure that the CSP can retain control of the end-to-end network and application performance and benefit from the distributed cloud.

It is more than the edge: Slicing and virtualization

At the same time, a distributed architecture gives CSPs the ability to manage application and traffic requirements in real time and dynamically respond to changing demand. For video-based services, this is critical, because the traffic load can be highly variable, but managed. For instance, CSPs can reduce video resolution to avoid congestion during an unpredictable traffic spike. The COVID-19 pandemic showed us how the ability to fine-tune traffic management may be highly valuable to ensure service continuity.

Virtualization helps manage workloads across the infrastructure in a flexible way. Depending on the traffic load, CSPs can shift processing across the cloud as needed, depending on the requirements of different applications or customers and the network resources availability and cost.

Network slicing is another powerful companion to the edge cloud. By managing traffic flows with different latency, throughput or reliability requirements differently, the value proposition and effectiveness of the distributed cloud can be amplified – and simplified as well. Network slicing

enables operators to decide which locations in the distributed cloud should support specific applications or services, based on what is available and on what the requirements are.

Building the foundation for future use cases

Video surveillance and analytics are a key learning ground for CSPs as they roll out 5G networks. Many 5G use cases rely on video uplink and downlink streams. The ability to hit the ground running by deploying the distributed cloud on an already established use case, such as video surveillance, will accelerate the move to other, emerging 5G use cases – AR/VR/XR, industrial automation and automated guided vehicle (AGV) applications, to name a few – with great long-term potential.

About Monica Paolini

Monica Paolini, PhD, is the founder and principal of Senza Fili. She is an expert in wireless technologies and has helped clients worldwide to understand new technologies and customer requirements, create and assess financial models, evaluate business plan opportunities, market their services and products, and estimate the market size and revenue opportunity of new and established wireless technologies. She frequently gives presentations at conferences, and writes reports, blog entries and articles on wireless technologies and services, covering end-to-end mobile networks, the operator, enterprise and IoT markets.

Nokia OYJ
Karaportti 3
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: SR2006044977EN (June) CID207624