INTELLIGENCE AT THE EDGE

GARY MARKS FROM OPENGEAR CONSIDERS HOW THE NETWORK EDGE IS EVOLVING WITH THE RISE OF THE INTERNET OF THINGS AND CENTRALISED COMPUTING

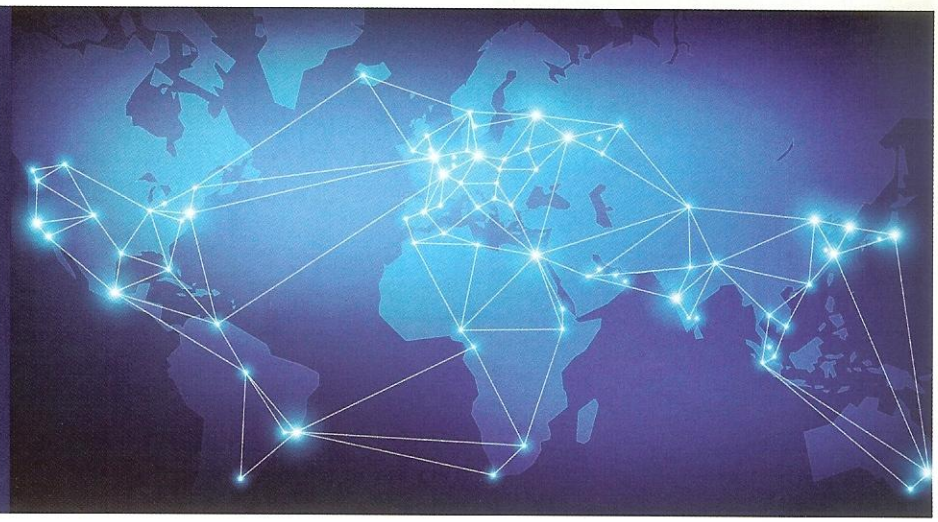

The current hype surrounding the Internet of Things (IoT) envisions a world where billions of devices at the network edge gather data and enact commands to create innovative and intelligent systems. From self-regulating traffic lights improving traffic flow, to smart homes able to adjust power and heating based on occupancy, the IoT nirvana offers a huge amount of potential for society. Yet a more substantive migration of technology to the edge has been quietly taking place over the last decade.

Examples abound. In financial services, bank branch offices have lots of screens but little compute resources running locally. Instead, thin clients communicate over secure networks back to centralised applications located within resilient data centres. Similar practices are common in retail and healthcare where pressure to provide increased cost benefit and improved security has led to a centralisation of IT within the core. Although the network edge has less compute, the criticality of local devices remains.

Now, local IT appliances such as routers, firewalls, edge caches, application traffic management switches and encryption equipment form a critical layer between the core and the edge. Yet these vital devices, often considered as appliances, are actually pretty dumb. Many are designed to do one job extremely well but don't deal well with unexpected changes or external issues.

To give an example, a number of devices within this category impacted by the recent Heartbleed security vulnerability had few ways to mitigate the issue due to limited configuration or even upgrade options. It's not just security issues. The devices on the edge are designed with the assumption that the IP network they rely on to communicate with the core will always be available. However, they lack intelligence to find alternative communication paths in the event of a core network outage.

In parallel with the rise of critical edge connectivity and core compute, many more organisations have added out-of-band management (OOB) to mitigate these points of failure. Yet, in an IoT universe where billions of devices are on the edge, OOB also needs to evolve. With the growth of mobile 3G/4G coverage, the first step is for OOB to allow the core to always connect with the edge even if the primary network is down. The evolution to cellular OOB is a fine example of how to maintain remote network infrastructure. The next stage is intelligence. OOB needs to be smart. Even with the rise of IoT, the hundreds of millions of network switches, routers and other network devices probably don't need to be upgraded to provide a conduit for IoT traffic.

However, the edge needs to be smart, with more intelligence. For example, a retail store with a number of devices connected by a local Wi-Fi access point may need to send data over the business broadband to access services or connect to applications residing within head office. If the broadband goes down, then the next generation of OOB needs the intelligence to initiate processes to fix, or at least mitigate, any issues. This could include cycling power to attempt to fix the fault or routing network traffic over an alternate 3G/4G cellular path.

Intelligence could also extend to potentially checking the settings of the homehub type device and maybe setting it to a default configuration to offer baseline connectivity. This level of automation is being developed within the remote management community and the first set of real world use cases are already emerging.

The levels of automation will vary depending on the scenario, but as the nature of computing and the network edge evolve, organisations evaluating their longer term compute strategy need to build resiliency extending from the data centre out to the edge. NC