Pipeline
Technology for Service Providers.
www.pipelinepub.com Volume 15, Issue 3

Intelligent File Transfer as part of Service Agility in decentralized IoT Networks

By:

Notice: Undefined index: issue in /home/http/pipelinepub.pubspoke.com/article.php on line 215

The growth in IoT is explosive, impressive—and unsustainable under current architectural approaches. CSP and DSP are moving away from a centralized network with challenges related to latency, network bandwidth, reliability and security. The new IoT business model requires decentralized networks and data lakes in order to support scalable service agility. As a result, service providers are partnering or building their own new data centers around the world.

As they do so, they are turning to software-centric architectures with the principles of Software-defined Networking (SDN) and Network Functions Virtualization (NFV) to unlock their constrained and disjointed network infrastructure. Globally, service providers are seeking to transition from a closed architecture with complex control and interoperability to a more open and efficient software-driven network and service architecture running on open, commodity platforms. This approach is driven by service providers' requirements for their networks to enable an agile service-driven environment that can dynamically respond to real-time subscriber demands.

The complexity facing our digital present and digital immediate future is huge. As part of our digital transformation journey, we will need to support millions of IoT devices in the near future, devices which in turn are producing tons of data, and it's desirable for that data to be ingested and processed at the edge of the network as close to the devices as possible. Some of that data— for instance, for autonomous vehicles—has to be sent back to centralized systems for further analysis as well as incorporated as part of overall system improvement in the respective big data.

The current disruption in the telecommunications industry and the respective accelerated digital transformation demand a new paradigm in supporting agile infrastructure and the creation of new services. Service providers today are challenged with moving big data— for example, a NFV image transfer from a centralized server to multi-access edge computing (MEC)—quickly and efficiently enough to meet business requirements. Currently most service providers leverage bandwidth utilization by FTP on high-speed lines over distance without acceleration. Furthermore those images and updates have to be conducted manually without a sophisticated distribution system. The reality is, though, that next-generation service agility demands new ways of thinking—and operating.

Not for distribution or reproduction.

Multi-access edge computing (MEC) emerged on the wireless industry stage several years ago. It is considered by many analysts and experts to be a technology as disruptive as anything that is being discussed today in virtualized networks—such as Network Virtualized Function and Software-defined Networking, C-RAN, and so forth. MEC is most likely going to help realize the promise of 5G. Simply put, MEC marries a radio with a data center at the edge of the network. The server component is a secure, virtualized platform which network owners can “open up” to third parties—including content providers and application developers, to name just a few. In doing so, the network owner allows content to be placed at the “edge”—that is, very close to the end consumer of that content. That content can be anything, including streaming video, augmented reality, location-based services, connected vehicle, and Internet of Things (IoT) applications. Digital transformation forces service providers to bring networks to the edge, where mission critical applications can be offered in real time. Applications such as AR, VR and autonomous cars require a response time far below the average for standard applications: for instance, under 10ms. In order to make this happen, mobile edge computing sites have to maintain the latest copies of applications in a mass market. The paradigm shift in new distributed networks foresees the usage of MEC. Now, we’re moving past bits and bytes to talk about how MEC is powering mobile network transformation to serve industries in innovative ways—for instance MEC uses a lot of NFV infrastructure to create a small cloud at the edge. And operators are starting to understand the economics that accompany MEC deployments as they understand how the technology works and what it offers.

Another way of supporting the access networks, primarily in fixed line, is fog computing. According to the OpenFog Consortium, fog computing is a system-level horizontal architecture that distributes resources and services of computing, storage, control, and networking anywhere along the continuum from cloud to things. Conceptually, it takes some—or all—of these resources down to the device level. And while MEC is intended for mobile networks, fog can include wireline networks.

MEC can be combined with other key technologies, like network slicing, to support 5G in new economic ways never before seen. Network slicing virtually carves up parts of the network to better allocate them for delivery of specific services. A slice meant to carry video might be tuned for high throughput, while a slice meant to power self-driving cars might be low-throughput but also extremely low latency. The idea is that the network stands ready to support a range of use cases in a way today’s networks currently cannot. Network slicing also suggests different cycles of NFV and container updates in the context of service life cycle management. By putting content and applications at the edge, the network owner can achieve operational and cost efficiencies while introducing new services, reducing network latency and— ultimately—improving the end consumer’s experience quality.

Early estimates suggest the number of towers in 5G will increase tenfold compared to the previous mobile versions. This proliferation would mean that—in the US alone—we would have to maintain, update, and upgrade over two million sites with the latest and greatest in applications and with NFV on a regular basis. In order to keep the MEF cells up and running, big data collected from cells have to be sent to central offices for further analysis. On the other hand, the agility of new services would require an instantaneous update of many sites, with additional challenges for multinational service providers that have numerous remote locations.

During the transfer of data between different locations and servers via the public Internet, leased lines or even dedicated lines, the transfer rate rarely reaches even half of the anticipated bandwidth. This adds additional cost in delivering new services, which can be especially problematic for connections with a large bandwidth or a high delay. In some cases, this kind of delay might be related to cross traffic or the fact that the bandwidth is used by another user or other applications. But we know for a fact, that the data transmission over private networks show no better results.

One of the most significant shifts is the evolution of data center-focused networks. Driven by the enormous expansion of cloud computing, the data center market is growing rapidly as business and enterprise users move intranet data service onto the internet. Between 2017 and 2021, the volume of worldwide traffic is expected to almost double, with 95 percent of that growth coming from the cloud or MEC, with modern high-performance networks to operate at speeds of 10, 40 and even 100 Gbps based. In case of large data transfer however, large bandwidth does not automatically

Not for distribution or reproduction.

correlate with faster and higher application traffic.

The Transmission Control Protocol is the de-facto standard for reliable data transmission across nearly all IP networks. Even with the increasing portion of multimedia traffic, about 90 percent of the bytes and packets transferred over the Internet are sent using TCP. To avoid congestion in a network, TCP uses a mechanism which reduces the so-called congestion window whenever it detects congestion. This leads to a short interruption of data transmission. According to this congestion avoidance mechanism, packet losses are considered as a sign of congestion. However, in real networks, there are several different reasons for packet losses. If a packet loss is not caused by a network congestion, reducing the transfer rate may not be effective, and it causes an unnecessary slowdown of data transmission. While there are a variety of techniques for users to fine-tune TCP to make it perform better, the main reason for poor performance still remains unchanged: the protocol itself.

In order to achieve a fast and efficient data transfer, a new transfer mechanism is needed, one superior to current TCP implementation. Over the last three years, Dexor has implemented a completely new patented protocol with the vision of creating a best-in-class protocol for fast and efficient data transfer. This protocol has proven to be superior to actual TCP implementation, as well as to proprietary data transfer solutions. It is called Reliable Multi Destination Transport protocol (RMDT). The keys to success for RMDT protocols are its algorithms for the analysis of the available bandwidth and very efficient data management in the protocol stack. Furthermore, the RMDT protocol can operate both in the classic domain of one-sender-to-one-recipient delivery, as well as deliver data with multi-gigabit speed to to multiple destinations simultaneously.

This process has propelled the project from the initial idea of overcoming the limitations of TCP towards the creation of a completely new, high-speed solution that fully utilizes the capacity of the available data bandwidth. To demonstrate the capabilities of RMDT, we have conducted a number of tests and compared the performance of several applications. The RMDT protocol can carry point-to-point as well as point-to-multipoint traffic, each with varying transmission speeds and MTU sizes.

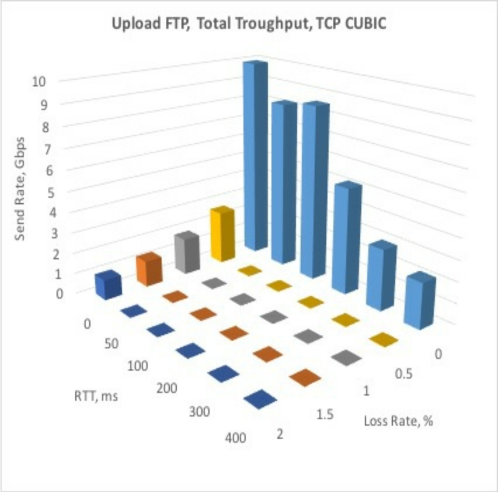

Figure 1 – FTP-based data transfer to multiple locations
(Click to enlarge image)

Not for distribution or reproduction.

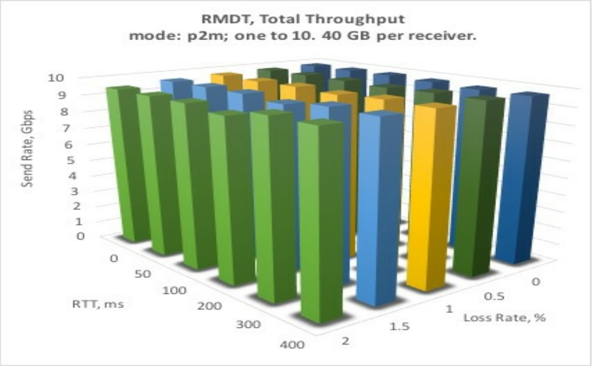

Figure 2 – Equally fast file transfer point to multiple locations using RMDT protocol
(Click to enlarge image)

The transmission in Fig. 1 is done using FTP-over-TCP (CUBIC) on a freshly installed Linux system out of the box. TCP CUBIC is an implementation of TCP with an optimized congestion control algorithm for high-bandwidth networks with high latency. In this case, the TCP has been tuned-up; however, the performance will grow only marginally. In Fig. 2, a point-to-multiple-point transfer with 10Gbps IP trunks is illustrated. Those tests were conducted in a simulation lab with delays up to 900ms in average (for long distance).

To verify the conducted tests for RMDT, sites were set up to showcase the power of the protocol. These included sites in Amsterdam, Vienna, London, Madrid, Denver and New York. These sites and their respective servers were connected via 10G line to the public Internet. The results have shown that the constant and guaranteed transfer speed decreased OPEX significantly and enabled service agility at scale for future MEC deployments. Big data applications can also collect a large amount of data, gathered in the IoT environment, and stream it back to centralized locations without saving the data on-prem, leading to higher efficiency.