Accelerate Time to Value

Ben Franklin originated the aphorism "*Time is Money*". This saying has stood the test of time. Today's economy, politics, and social order are based on the dispersal of information. It used to take days and weeks to relay information across distances but with the advent of technology, the time to connectedness has shrunk to hours and minutes. The new demand is the "speed of relevance" based on Access Time Objective (ATO). Often, the older the data is, even minutes and seconds, the less value can be extracted from it. It's similar to the expiration date on milk. If the expiration date is in two days, the value of the milk will be higher than if the expiration date is today, or worse, yesterday.

The IT industry has been focused on reducing time to value in every way, except one. The adoption of GPU processors has accelerated many analytical processes. We have seen faster CPUs, faster storage media, and faster networks. The one variable that has seen little to no progress is latency. Of course, the obvious reason is, latency is physics. We also haven't figured out how to travel through wormholes. TCP, a protocol designed for local traffic, has not adapted well to WAN transport. A unique flow traveling over TCP with latency above 50ms, will see ~10% of bandwidth utilization. We can send multiple flows simultaneously to fill the pipe, but data will still arrive at the same rate because the effects of latency on each flow are constant. So, we have developed new approaches such as using UDP to transport data, creating copies of data stored local to the application, using WAN optimization, which leverages spoofing protocol chatter, deduplication, and caching to reduce traffic over WAN. The improvements these techniques offer do very little to address the need for "speed of relevance".

Let's restate the challenge at hand: we can't eliminate latency due to the laws of physics. TCP was not designed for WAN traffic and UDP offers to guarantees that data reached its destination; caching and WAN optimization offers minor improvements but at a significant increase in cost. Time is money – missing ATO results in loss of opportunity, revenues, equity, and value.

Vcinity has set out to address the speed of relevance and access time objectives; not how much faster I can move data or store it near applications, but how can I access it. In our presentation at TechFieldDay #TFD23, Vcinity team will demonstrate how Vcinity Data Access Platform™ (VDAP) enables applications to process data anywhere over any WAN (Microwave, IP, Satellite, 5G), such as video rendering, explain our approach to eliminating the effects of latency, and share use cases where VDAP has made dreams come true. This doesn't drop as latency reaches 400ms, 500ms, or 600ms. This enables applications to access and process data as though data and compute were collocated on a LAN. It is not magic. Visit us on LinkedIn, vcinity.io, or during #TFD23 on April 21st and become a believer.