Pipeline
Technology for Service Providers.
www.pipelinepub.com Volume 13, Issue 8

CEM, Orchestration and Big Data

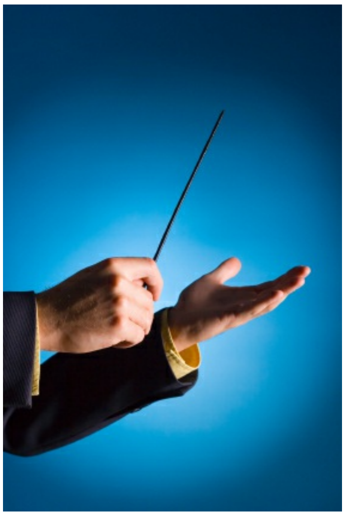

By: Mark Cummings

There are two approaches to Customer Experience Management getting a lot of attention today. They are End-to-End Orchestration (EEO) and Big Data. Some argue for one over the other as a single complete solution. While, in fact, they both have an important and complementary role to play. Similarly, there are differences of opinion about fully distributed vs. fully centralized solutions. Here again, the correct answer is not "either or" but "both and". To better understand these complementary relationships we will first briefly explore the relationship between Customer Experience Management (Quality of Experience or QoE) and network Quality of Service Management (QoS). Then we will examine real-time response and pattern mining and how they can work together. In this context, we will look at topology and examine how distributed and centralized approaches complement each other and are driven by external factors. This discussion is a first simplified high-level view.

The Relationship of QoE and Network QoS

QoE management focuses on keeping underlying communications systems (whether virtual – VNF or physical PNF) functioning at a certain level of operational quality. CEM focuses on keeping a service as experienced by an end user functioning at a certain level of customer perceived quality. It is possible for Quality of Experience (QoE) to be high on a set of networks that have low Quality of Service (QoS). In this situation, multiple networks or sub networks, for simplicity let's call them paths, are employed. As long as the degradation or failure rates of the various paths are statistically independent, and there is an ability to very quickly switch between them, high QoE can be provided even if the QoS of the underlying paths are low.

Of course duplicating or triplicating resources lowers overall efficiency and raises cost. In the short-term, it may be that the cost of poor QoE (customer churn, penalty fees, etc.) may be higher than the cost of lower efficiency. But in the longer term there is generally an incentive to improve efficiency and lower costs.

Orchestration and Real Time Response

Properly constructed and deployed End-to-End orchestration (EEO) insures high QoE. EEO is the way that the various paths are switched in and out quickly enough at the QoE level. To be effective at QoE management, the end-to-end capability is very important because the different paths may involve different domains (RAN, WIFI, Wireline, Fiber, Core, etc.), different vendors, and different administrative units.

Response time is very crucial for QoE. The objective of QoE Management is to provide something akin to traditional fault tolerance with switch-over times fast enough that the end user is either unaware or unconcerned. This goes beyond the simple satisfaction of a Service Level Agreement (SLA). SLA's are in essence a worst case performance hurdle. QoE's target is to achieve best case performance. Some people talk of the response time requirement for QoE as "real time". The concept of real time performance comes from computer operating systems and has to do with such things as non-maskable interrupts. The problem with that, is that there is no good metric for real time in a complex Telco network.

Response time does not only depend on the EEO, it also is a function of how fast the interfaces to the underlying equipment can respond. The kind of adjustments that the EEO makes to manage QoE most closely resemble provisioning (the set up of network components). Because legacy systems were designed in a time of manual interfaces where provisioning took hours, days, and sometimes months; many of these systems have very slow response time interfaces. So, while the EEO can respond in less than a second, the underlying system components may not be able to.

As VNF's (Virtual Network Functions) are rolled out, they tend to implement the legacy interfaces with the legacy sluggishness. Therefore, we use a grading system instead: A for less than a second; B for less than a minute; C for less than 10 minutes; and D for less than an hour. As time goes on, and the requirement for rapid response becomes clear, response times may increase and our grading system ratchet up.

The role of End-to-End Orchestration in QoE

LTE Advanced has a function called SON that does load balancing between basestations (eNodeB's). In metropolitan areas it is common for a UE (User Equipment - handset, etc.) to be in range of a number of basestations. Ten is a rough average. LTE SON works to distribute UE's in such a way as to provide the best QoE for all active users. Unfortunately, it makes its decisions only based on the RAN (Radio Access Network) capacity of each basestation in a neighborhood. This capacity is set in initial provisioning and is often independent of the underlying capacity of the basestation. This load balancing is blind to backhaul capacity. Backhaul capacity can be impaired by component failure, traffic congestion from other traffic sources, weather, or other factors. A rough average is that at any given time, 1 percent of the basestations have impaired backhaul. When that happens, users experience dropped calls, delayed or dropped texts, delayed emails, poor or no response to web searches, and so on. Users, respond by trying multiple times, which increases the overall traffic in the neighborhood creating congestion on the other basestations. The QoE of the whole neighborhood is affected. This can mean that 10 percent of the neighborhoods have poor QoE.

EEO manages QoE by first detecting that the backhaul impairment has occurred and then reducing the RAN capacity of the affected basestation by moving part of its capacity to its neighbors. This assures the UE's assigned to the affected basestation good QoE, stops all the retries, and maximizes the usefulness of the capacity (combined RAN and Backhaul) that exists in the neighborhood. When the backhaul capacity of the affected basestation has been returned to normal, the EEO moves RAN capacity back to the affected basestation.

In a 100 million subscriber network with one million macro basestations, there are 10,000 basestations with impaired backhaul (with average neighborhoods of 10 basestations) affecting 100,000 basestations. Each LTE basestation has 6,000 software settable parameters many of which generate large volume data streams. Then, these data streams must be combined with the data streams from the many switches and routers that make up the backhaul network.

In many cases the backhaul network also provides connectivity to central cite processing systems. So, a problem in backhaul may also make it difficult for the fault information to reach a central site in a timely fashion. This makes it very difficult to do EEO from a single central site system. Better QoE can be provided by a distributed system working at the edges. As the information needed exists at the edge and all required changes can be made at the edge. There may be a need for some central coordination, but the most efficient way to provide this kind of QoE management is at the edge.

The Role of Big Data

At the same time there is a need for aggregating data to find patterns that exist across large portions of the network. This is the role of Big Data. Time constraints are relaxed and the focus is on pattern recognition. Once an actionable pattern has been identified, then it can be passed to the EEO system for short-term execution, and to network planning, and other systems for longterm.

For example, let's look at a variation of the use case described above. If, instead of a backhaul failure, there is a failure that affects a basestation itself, the EEO system can also move RAN capacity to other basestations in the neighborhood. This improves EEO in the short-term. At the same time the Big Data system is looking at failure rates of components and it discovers that, on average, a particular board in a particular basestation model fails in 18 months (a Telco tells me that such a pattern has been found). So, now both the EEO system and the network planning system track when these particular boards are put in service. Instead of waiting for them to fail, they are now replaced when they have been in service for 17 months. With NFV, as Telco networks begin more and more to resemble Clouds, anticipation based on actionable patterns may find more places to add value.

However, once having found such an actionable pattern and created algorithms based on it, the work can't stop there. There has to be ongoing work to check and make sure that the patterns have persisted. For example, has the board vendor made a software or hardware upgrade that extends the mean time before failure? If so, the Big Data system by itself will not see the new pattern, because all of those boards have been removed. So there has to be some form of ongoing test. In this use case, running an ongoing test of a sample of removed boards to see it the failure rate has increased, while having the EEO system report in service outages to see if it has decreased.

Thus it can be seen that there is a symbiotic relationship between EEO and Big Data in QoE Management. Neither is sufficient by itself to assure the highest possible QoE, but working together they can.