Leveraging machine learning to eliminate backhaul bottlenecks in 5G networks

white paper

Leveraging machine learning to eliminate backhaul bottlenecks in 5G networks

white paper

Anssi Tauriainen
Director of service assurance and analytics business unit, EXFO

Studies have shown that many service degradations and outage issues are caused by connectivity issues in the backhaul network rather than with the radio network.

The backhaul challenge

Mobile networks are becoming increasingly complex and discovering the root cause of service and network issues is not always obvious. In fact, studies have shown that many service degradations and outage issues are caused by connectivity issues in the backhaul network rather than with the radio network.

Connectivity in the backhaul network can be impacted by any number of issues. Congestion is, perhaps, the most common issue, but even that can be difficult to pin down since the congestion could be caused organically through increased traffic volumes, uptake in new service offerings or, inadvertently, through operator initiated actions, such as updated network routing configurations or changes in policy or firewall settings. Regardless, finding and fixing the root cause quickly is essential to meeting service level objectives.

In today's large backhaul networks, studies have shown that up to 1% of the cell sites are suffering backhaul issues, and while that might not seem like a lot, the impact on customers being served by those sites can be severe, resulting in customer satisfaction issues and potentially loss of revenue.

And the problem will only get worse as technologies such as Fiber-to-the-antenna (FTTA) and Centralized Radio Access Network (C-RAN), driven by the explosive growth expected from 5G, increases fronthaul capacity by up to 10 times the current LTE levels. In fact, a recent Light Reading survey indicated that 33% of CSPs expect backhaul performance to be the biggest obstacle to 5G rollout.

100% visibility with active testing

So how do you detect and isolate connectivity related issues, especially in a virtual or hybrid network where the dynamic nature of the abstracted service layer means there is no longer a 1:1 relationship between service topology and network topology? For this, active testing is ideally suited. Active testing involves generating and sending synthetic test traffic between the service endpoints and using that traffic to measure the performance of the service.

There are 2 key benefits for using active testing. First, because the test traffic follows the same path as the service traffic, the test accurately reflects service performance, in real-time. Second, in a virtual or hybrid network, where service routing can change, the test traffic automatically remains with the service traffic meaning no need to reconfigure test configurations and no gaps in test data.

Another key benefit of active testing is its' inherent scalability, especially in virtual or hybrid networks. By including an active test VNF in every service chain, 100% visibility of the services in the network can be easily achieved. Active test environments can easily support hundreds of thousands of endpoints and tens of millions of tests per day, covering all aspects of the service stack, from layer 2-7.

A typical test set up for mobile backhaul networks is presented in figure 1.

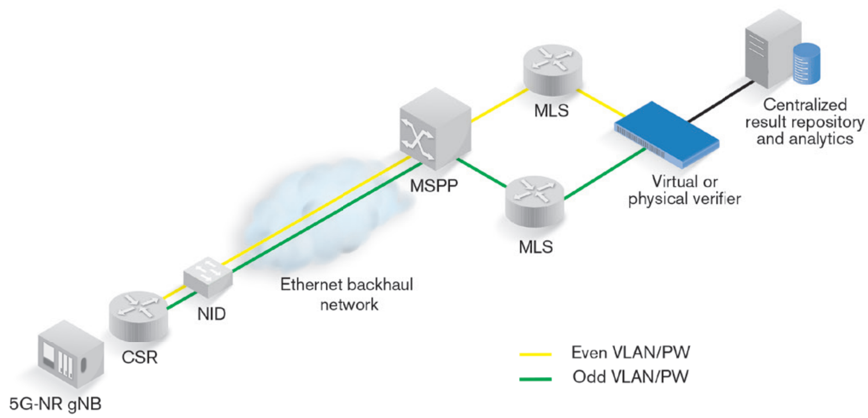

Figure 1. Typical active test setup for mobile backhaul in 5G phase 1 network

According to several market studies, active testing will become an essential part of SDN/NFV networks to ensure the proper functionality of newly created virtual services and functions.

Today's solutions must be able to automatically learn the behavior of the network and discover the various events that can lead to service degradation and outages.

Machine learning to address scale and automation

Traditionally, service assurance and analytics solutions for mobile backhaul are based on heuristics, such as rule based alarming, to notify the operators when pre-defined thresholds have been exceeded. And while this approach can be done in real-time, it is still reactive and triggers action only after an issue has happened. It also relies entirely on the configuration of pre-defined thresholds, meaning it's nearly impossible to cover all possible scenarios, due to the complexity of the network and volume of the data. Additionally, it is virtually impossible to distinguish between different failure modes to provide meaningful root cause analysis support.

Today's solutions must be able to automatically learn the behavior of the network and discover the various events that can lead to service degradation and outages. Doing this reduces manual, time consuming and often error prone work by isolating and resolving issues automatically and, in more complex scenarios, may even provide sufficient warning to operations personnel to react before issues escalate.

There are two main root cause scenarios to be considered: sudden short-term outages and longer outages.

In the first scenario, the outage usually starts without any early warning signs and is causes temporary service interruptions. Typical root causes include events such as a node reset or restart in some part of the network. Typically, the network can react automatically, and situation is recovered relatively fast without any intervention.

KPIs are typically network oriented and, therefore, may have no direct correlation to end user experience.

The second scenario is more severe and usually requires human intervention for recovery. Situations like this are caused by non-recoverable failure in the network, such as software faults or misconfiguration, hardware failure, or capacity exhaustion, either temporary or permanent.

In these cases, there are often early warning signs, such as degradation of the performance in some part of the mobile backhaul network. This degradation causes problems, such as congestion in adjacent parts of the network, leading eventually to degraded service levels at individual cell sites connected to that part of the network.

The first, obvious requirement for a self-learning system is to be able to predict an outage, but it is equally important to distinguish the type of outage. For example, in the sudden short-term outage scenario, it is sufficient to identify that an outage has happened and recovered automatically.

However, in second scenario, it is essential to be able to predict the outage and alert the operations personnel as soon as possible. For this reason, self-learning system must be able to analyze the events that took place prior to the outage to detect failure patterns.

Based on these scenarios, the self-learning solution must be able to perform following sub-tasks:

1. Identifying an outage
2. Finding causal factors for the degradation
3. Classify the outage
4. Execute right corrective actions

Two approaches to handling data

Backhaul assurance and monitoring systems utilize various protocols in testing the network connectivity and performance - the most common being RFC 5357, Two Way Active Test Protocol (TWAMP). In TWAMP test packets are sent between two end points, as illustrated in figure 1 above, and various measurement results recorded key performance indicators (KPIs) such as frame loss or frame delay variation (also known as jitter). TWAMP produces tens of KPIs per test end-point, and test packets are being sent and received all the time, so in practice the active assurance system produces a constant, very large data set for analysis purposes. This is good news for machine learning algorithms, since large data sets are needed for producing accurate results. But it also means that the infrastructure must be capable of handling, potentially, terabytes of data in real-time.

In machine learning systems, there are two commonly used approaches that have been successfully implemented, each brings its' own set of strength and weaknesses.

Unsupervised anomaly detection – where the machine learning system is fed with unlabeled data, and the algorithm learns the anomalies within the data mass. Such algorithms can discover relationships within the data which would be otherwise undetectable.

With right preparations and data, a machine learning system is capable of predicting backhaul incidents with a high rate of accuracy.

The obvious benefit of this approach is that setup is easier, and less data preparation is needed. For example, there is no need for data labelling, which can be a tedious task, especially when talking about complex networks and data. Another benefit of this approach is that no separate data set is needed for training. These principles are illustrated in figure 2 below.

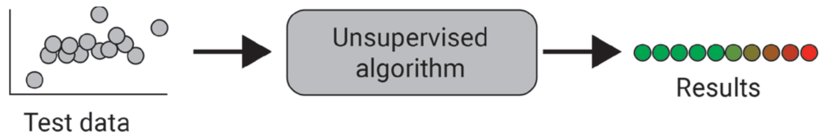

Figure 2a. Unsupervised anomaly detection

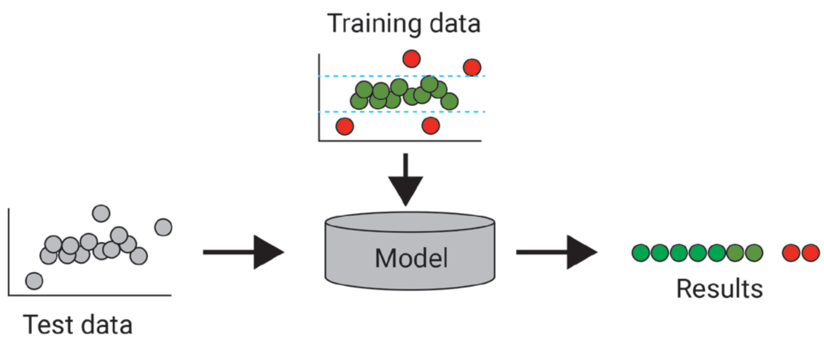

Figure 2b. Supervised anomaly detection

Unsupervised learning algorithms are especially useful when there are several types of different source data available. For example, finding an unexpected correlation between network performance data, network topology, weather reports and trouble tickets is possible. It should be noted however that the method produces false positives and new discoveries need to be manually evaluated to confirm the relevant ones. For example, the correlation with a weather report is relevant when part of the backhaul network is implemented as microwave link, but should have no correlation to fiber backhaul.

Supervised failure prediction is effective when data can be labelled as “normal” and “abnormal”. It requires more preparation but once done, the algorithm can automatically distinguish the scenarios and find the issues contributing to “abnormal” behavior of the network.

The labelling of desired “abnormal” behavior in backhaul networks may not be as straightforward as it seems. KPIs are typically network oriented and, therefore, may have no direct correlation to end user experience. For example, while 100% frame loss on single path indicates that there is a total interruption in data transfer between two end points, it is still possible that a cell site has a backup connection which can handle end user traffic without any interruption during primary connection outage. Including external data, that measures the end user service level being experienced, may be needed. For example, Key Quality Indicators (KQI) such as drop call ratio, download speed or network latency measured from the end user device would be ideal for labelling purposes.

Supervised algorithms are capable of learning to predict outages in large data sets with reasonable accuracy and high success rates. It is possible to increase the accuracy but only at the cost of success rate, since the algorithm starts to create an increased number of false-positive predictions.

When investigating the predicted incidents, analysis of the causal factors (early warning signs) can provide an additional means for the automatic classification of incidents. For example, a combination of high frame loss and frame delay variation usually indicates a gradual exhaustion of capacity. Or a cell site experiencing higher frame loss than neighboring cells is probably experiencing equipment failure.

Conclusions

Mobile backhaul networks will likely be the next bottleneck when networks evolve toward 5G and C-RAN infrastructure as capacity and performance will not be sufficient to serve the new services and increased bandwidth in fronthaul and RAN. By leveraging advances in machine learning for automated assurance of the backhaul network, operators can not only predict outages and service degradation but also forecast capacity bottlenecks.

Supervised learning algorithms with proprietary labelling are more suitable for assurance and forecasting purposes than unsupervised learning algorithms. However, this requires extra effort on the part of the operator to label the data and additional data sources may be required for identifying abnormal network behavior.

With the right preparations and data, a machine learning system is capable of predicting backhaul incidents with a high rate of accuracy. It is also possible to teach the algorithm to classify the types of incidents, which in turn shortens the time spent on troubleshooting and eventually helps the automation of corrective actions.