For network slicing, service assurance can't be an afterthought

By: Ross Cassan

It's almost become a telecom tradition for operators to push a great new offering into the market, only to eventually struggle with maintaining quality and delivering a great customer experience as the service matures. The stumbling blocks were always going to be there, but because service assurance wasn't planned from the start, operators kept themselves in the dark. By the time they were able to turn on the light with a proper service assurance strategy, upwards of a year or more may have passed. And historically, operators have paid the price in frustrated consumers.

Things will be different with network slicing on 5G and all the mission-critical services it will be expected to support. The stakes will be much higher, necessitating that service assurance be baked-in from the start.

Operators are planning now how 5G networks will be sliced to meet customized performance needs in an expanding range of verticals. While approaches will vary by region, stringent demand for service levels unlike anything they've had to provide before will be a constant. Missing the mark could mean regulatory fines and SLA violations that eat significantly into revenues. Even more serious, however, is the potential fallout from issues that affect emergency services, public safety, smart cities or autonomous cars, just to name a few.

Guaranteeing great experiences in networks that constantly change

Spirent has been working closely with device makers, network equipment manufacturers and operators on early 5G testing. Working across the full ecosystem has given us a unique vantage point for triangulating expected issues. We see immediate areas of concern at this stage:

- The dynamic nature of network slicing insists the network be continuously reconfigured on the fly to meet the ultra-low latency, high-bandwidth demands of specific applications. Every time this happens, the network is being changed in significant ways. It's a process that will happen over and over in 5G networks.

- The core network, which used to sit in stable, centralized data centers, now may be pushed to the network edge to deliver enhanced performance. Complicating matters, the control plane and user plane are now separate and independently configured. The distribution of network functions and control user plane separation (CUPS) are radical changes, introducing a new layer of complexity that inevitably results in performance issues and configuration challenges.

Previously, turning up new network infrastructure was a manual practice that could take weeks or months. Lab-based validation was often the long pole in the tent, requiring the use of specialized physical test probes and manual processes. In 5G networks, testing will need to be dynamically triggered, occurring in seconds or minutes. That means it must become software-based, take place in the operational network and be part of an automated workflow. In other words, the testing conducted to assure services needs to be part of how the service is designed.

We know that some infrastructure providers will shout from the rooftops that they plan to handle assurance and QoE requirements in the network gear they deliver. But the simple fact is if these built-in telemetry capabilities aren't testing the end-to-end network slice performance—right before it goes live for the customer—providers can't be sure critical SLAs will be met. The industry has not quite woken up to this fact yet and our cutting-edge work in the lab with 5G tells us there will absolutely be issues, from missed latency targets and lack of resources, to misconfigurations and beyond. These potential issues can only be solved with automated, proactive testing that can detect fleeting flaws as they happen, isolate responsible network functions and provide actionable guidance to network management and orchestration layers.

Virtualized assurance gives visibility of the network edge

Let's say a customer needs an ultra-low latency network slice for a particular application like augmented reality or connected vehicle guidance. 5G makes it possible to dynamically create network slices where core network functions are pushed to multi-access edge computing (MEC) nodes, significantly reducing service latency.

But how do you ensure the new slice will work after it is instantiated? 5G service assurance will need to deploy virtualized test functions (called Virtual Test Agents) to MEC nodes. These VTAs will then pretend to be users or other parts of the network and perform tests on distributed network functions to make sure they will deliver the low latency or other critical performance needs of the end user. Once the service is up and running, the entire network slice can be monitored to ensure end-to-end performance SLAs are continuously met.

5G designs must incorporate automated, virtualized assurance from the start

The dynamic, distributed nature of 5G makes traditional manual, "bolted-on" service assurance useless. By the time a problem is detected with a traditional approach, the 5G network configuration will already have changed. What's more, traditional approaches simply don't give visibility out to the network edge – a key need for 5G distributed architectures. But a network designed from the ground up with automated, virtualized service assurance can dynamically assure new network slices right out to the edge of the network. This approach also overcomes a key limitation of telemetry, providing visibility of the end-to-end performance of 5G services from the customer's perspective.

5G needs automated assurance, but many providers are still working out their broader operational automation strategies. Key enablers of operational automation such as ONAP (Open Network Automation Platform) and cloud-native networks are highly complex, and most providers aren't ready for comprehensive, network-wide deployments yet. So how do we automate 5G assurance when all the automation ingredients aren't quite ready?

5G automation use cases for early deployments

We recommend getting started with a basic set of automation use cases that lay a foundation for broader automation initiatives, but don't require all elements to be in place:

Service turn-up and activation – *Automated validation of new services, network slices and network functions before they go live.* To automate these workflows, virtualized service assurance systems must feature a cloud-native design that allows rapid integration with network management and back office systems. 5G VTAs need the ability to accurately mimic real-world network functions and devices, so they can test new functions, slices and services in advance of actual users. Applying this type of automation for 4G has been proven to accelerate turn-up and activation by a factor of ten for tier one providers. We expect a heightened importance in 5G since manual techniques are no longer viable.

Change management – *Automated validation of all network changes and upgrades to ensure new services and features meet performance targets without degrading the quality of existing services.* In 5G, this will require VTAs that can surround a new or upgraded network function, emulate the rest of the network and verify performance before the switch to go live is flipped. One key goal of this automation is to avoid "micro" network or service outages where the provider is unaware an issue has occurred until a customer complains. When a tier one mobile provider applied this approach to 3G and 4G networks, undetected issues were reduced by 85 percent. Because one of the key selling points of 5G network slicing is unprecedented latency and speed performance, we expect providers to require even lower levels of undetected issues, particularly for critical applications.

Service quality monitoring and fault isolation – *The ability to proactively and continuously test service quality and performance across the end-to-end network and automatically isolate issues to specific network functions or infrastructure.* In 5G networks, we will see an explosion of applications that send traffic only at specific times or when certain conditions are met (e.g., alarm systems, IoT sensors, business applications, etc.). Virtualized service assurance systems need to be able to proactively test these services and underlying network slices to ensure they work when called upon. The systems also must be able to test any part of the network, including the radio interface, to rapidly isolate issues. Automation of these processes for 4G is already saving tier one providers millions of dollars annually in the form of reduced SLA penalties and operational troubleshooting expenses. With the increased complexity of 5G network slicing on top of legacy networks and new IoT and business use cases, we expect an even greater need for automation in this area.

5G is coming faster than most expected, and in this race, not a single corner can be cut—especially when it comes to network slicing. Next-gen mobile technologies will set the stage for operators to make bold promises to their customers. Service assurance is what will mean they can keep them.