TESTING NON-TERRESTRIAL NETWORKS

Introduction

A non-terrestrial network (NTN) is a wireless network consisting of nodes that aren't physically connected to the ground. The nodes can be uncrewed aircraft (drones) 5 to 30 miles above the surface or satellites at various altitudes, from low earth orbit (LEO) to medium earth orbit (MEO) to geostationary earth orbit (GEO), or even highly elliptical orbit (HEO).

This detached operation means that they can provide coverage to remote areas where it would be handy to monitor or control things, but impractical, or even impossible, to run cables or build wireless transmission towers.

NTNs have a wide range of applications.

- **Internet of Things.** Monitor goods and vehicles across global shipping routes; monitor crop health, irrigation systems, and livestock; manage offshore oil rigs using real-time data on equipment status and environmental conditions.
- **Direct to Cell.** Make calls through the satellite network using only your smartphone.
- **Maritime.** Track vessels and monitor environments in open ocean or coastal areas with limited connectivity.
- **Commercial Aerospace.** Characterize and optimize the performance and resilience of safety-critical avionic systems as well as infotainment systems, increasing speed-to-market.
- **Industrial.** Monitor and track remote assets, monitor environmental factors, and manage facilities.
- **Defense.** Create tactical advantage by enabling effective command and control through enhanced real-time situational awareness in all theaters, including facility security, managing environmental systems, and military logistics systems.
- **Automotive.** Provide advanced infotainment, collision avoidance, and driver-assistance features in transit.
- **Public Safety.** Support reliable and continuous communication with first responders regardless of location.

Anatomy of a Non-Terrestrial Network

There are a few simple concepts required to grasp the basic structure and functionality of NTNs.

Infrastructure

Low earth orbit (LEO) satellites orbit the Earth at an altitude between 300 miles or 483 kilometers (orbital period of about 88 minutes), and 1300 miles or 2093 kilometers (orbital period of about 127 minutes).

Medium earth orbit (MEO) satellites orbit the Earth at an altitude between 1,243 miles or 2,000 kilometers (orbital period of about 2 hours) to 22,236 miles or 45,385 kilometers (orbital period of about 12 hours).

Geosynchronous earth orbit (GEO) satellites orbit the earth at an altitude of 22,236 miles above the equator, with an orbital period matching the earth's rotational period, which means that from the ground they appear motionless at a fixed position in the sky.

Links

Service link. From a ground-based user to a satellite.

Feeder link. From a ground-based gateway to a satellite.

Inter-satellite link. From one satellite to another.

Technology

User plane. The data path that carries calls, messages, and internet traffic.

Control plane. The path used to manage network operations, including routing data, user authentication, and resource allocation.

Synchronization. The process of aligning the timing and frequency of devices communicating within the NTN. Synchronization compensates for the propagation delays and Doppler shifts found in space-based communication environments.

Satellite handover. Unlike GEO satellites, which maintain a constant distance between each other and earth stations, a LEO or MEO satellite orbit period varies between 90 minutes and 10 hours. Consequently, the distance between any given satellite and a ground station constantly changes. As a satellite moves outside the range of a base station and another comes within range, the system performs a handover from one satellite to another to maintain unbroken communications.

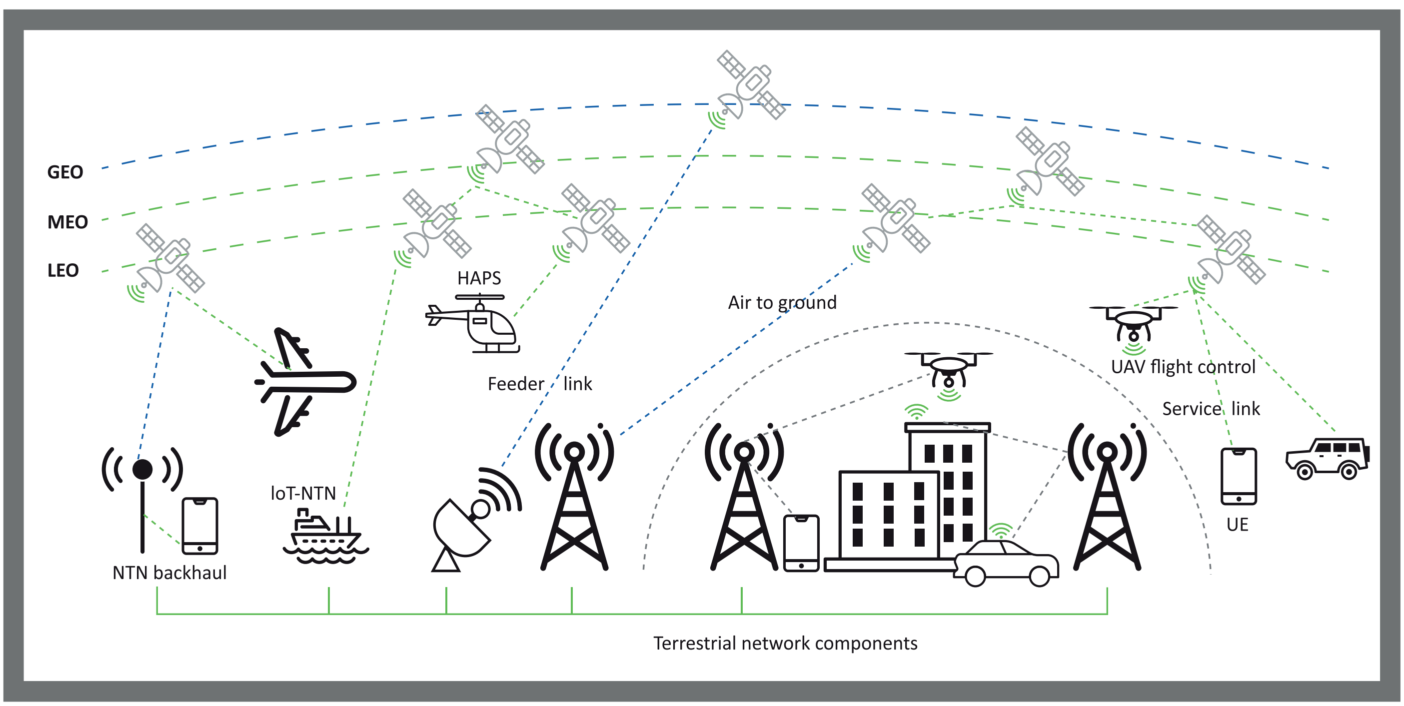

Overview of Non-Terrestrial Networks (NTN)

Challenges and Complexities of NTNs

NTNs inherit all the challenges and complexities of terrestrial networks, plus a few extras. These challenges include the vastly increased distance between network nodes, variable weather conditions, adjacent satellite interference, and handovers prompted by satellite movement. These dynamic network conditions can affect link integrity and performance quality.

Because of the amount of time it takes a signal to travel between a ground station and a satellite, latency values are higher in NTNs than in a terrestrial network, such as a cable or fiber internet link. These greater distances can account for up to a 20x increase in latency. In addition, due to congestion and atmospheric conditions, NTN latency can also vary widely.

Despite these challenges, customer expectations for connection reliability and application performance remain high.

Latency values are higher in NTNs than in a terrestrial network.

Testing Considerations

Solutions gain traction in the market for two reasons:

- They solve a problem quickly, effectively, and efficiently and they stand the test of time.
- There's only one way to validate that your solution has these characteristics.

Test it under the same conditions the customer will use it.

The problem is that you can't test the functionality and performance of a pre-production solution on a production network. Not only will it disrupt the users of the network, it but it will also fail to produce meaningful results. A production network doesn't have the three essential elements you need in a test bed.

It must be accurate, reliable, and repeatable.

Accurate	The test results should closely reflect the actual network behavior, minimizing errors and providing reliable data about network performance.
Repeatable	If the same test is run multiple times under identical conditions, it should produce consistent results, allowing for comparison and identification of trends.
Reliable	The testing process should be dependable, consistently delivering accurate results even when conducted by different individuals or at different times.

In addition to these three essential elements, a test bed should be able to recreate the many and varied conditions that can occur out in the wild. Failure to test, or failure to test rigorously under real-world conditions, can lead to system outages, increased maintenance costs, and customer churn.

Pre-production Testing

All testing begins in the lab, a place where you can control all the variables to isolate and stress a specific aspect of your solution. It starts with characterizing the performance of a range of elements, beginning with validating the most fundamental functions, such as forwarding bits properly, and moving up incrementally to validate higher levels of functionality and complexity.

Type of Test	Description
System integrity testing	The baseline of testing is the bit error ratio test (BERT), which assesses device performance by sending a known data pattern through the system and then analyzing how many errors occurred during transmission.
Functional testing	This suite of tests assesses the features of the device, such as alarms, fail-over, security, and such.
Latency characterization	Latency tests measure and analyze the time it takes to send or forward a message.
System performance testing	System performance tests evaluate how well a complete system performs under varying workloads. It analyzes key metrics, including bandwidth, latency, throughput, packet loss, and response time, to identify potential bottlenecks and optimize data transmission across the network.

The goal of this suite of test methodologies is real-world testing.

Real-world testing subjects the system-under-test to the same conditions that it will experience in the real world, a.k.a. the production network.

Post-production Testing

Once your solution is ready to go out into the real world, you will need to verify that it can keep up with what the real world will throw at it. To do that, you need a test bed that can:

- Measure end-to-end latency with high precision.
- Capture packets with high precision timestamps for protocol analysis, accurate event-timing analysis, and troubleshooting.
- Reproduce real-world network conditions in the lab to address specific issues happening in the wild.
- Playback captured production traffic to characterize a found issue and identify the root cause.

The proper test bed will ultimately reduce mean-time-to-repair, thus avoiding additional cost and penalties.

Additional Test Scenarios

Going beyond pre-production and post-production testing, there are other scenarios that bring value to your organization.

1

PROOF OF CONCEPT (PoC)

Once you have determined that your system can match or exceed the performance of the alternatives, you need to communicate the extraordinary value of your solution to your customers and investors. That is the function of a proof-of-concept demonstration.

A PoC test isn't performed in a sterile lab network under ideal conditions, but in a realistic emulated environment to demonstrate that your solution can adjust to punishing, real world conditions and still outperform other solutions.

You can show, not tell, that your system can optimize network architecture choices, avoid overprovisioning, reduce time to market, and much more.

2

PERFORMANCE CHARACTERIZATION

When you have designed a system for a specific environment and application, it's important to establish the performance characteristics of the solution to verify that it will address your customer's needs with plenty of margin. That is the purpose of a performance characterization test.

Performance characterization presents your solution with both normal operating conditions and also worst-case scenarios subjecting it to real-world dynamic conditions that are present in Satcom or NTNs.

Validate throughput rates. Characterize application response time. Verify system reliability against real-world conditions. Demonstrate that fail-over is triggered as expected. In short, accurately demonstrate the application or service performance they will experience before actually deploying anything.

By characterizing real-world performance, you can optimize applications and services, tune protocols, and properly right-size the system and network designs.

Choosing the right Testing Solutions

There are several metrics that are used to identify the best testing solution for NTN systems.

First and foremost, a test must be repeatable. This factor should be non-negotiable. In other words, you must be able to use the test results to identify issues to resolve, take steps based on the results to resolve the issues, and then repeat that same test to verify that all issues have been resolved. Anything else will leave you playing whack-a-mole or chasing ghost symptoms that come and go.

Repeatability can be accomplished only with high performance test system that is capable of precise measurement.

High performance means a test system performance greater than the system under test (SUT). For example, if the SUT can handle greater throughput than the test system, then the test will not reveal the true capabilities of the SUT.

Likewise, if the test system is not capable of sufficient precision in measurements such as latency or packet/event timestamping, then the measurements will be varied, a.k.a noisy, and the results will be impossible to interpret in a consistent and repeatable manner.

The above issues ultimately can cause the system under test to behave differently from test run to test run. The result is that the value of your testing is fundamentally lost. And in fact this can cause significant disruption in time and cost, making it difficult to figure out what is truly going on, or perhaps worse, wasting time chasing problems that don't even exist within the system under test.

Second, a test system must be flexible. It should cover a broad range of use cases and interfaces. The last thing you need in your crowded test lab is having to make room for multiple test platforms to support various interfaces or use cases.

And finally, it must be simple to use. It should minimize the time-to-know, the elapsed time between having a question about the performance of your system and getting the answer. And it should be accessible to all members of your team, regardless of their experience or expertise.

Take your readiness to the next level

With satellite technology now actively being deployed in a wide range of applications, smarter test methodologies are imperative.

Specifically, if you are not considering the effects of delay, limited bandwidth, and other network impairments early in your development and testing, you are likely to encounter irregular performance and failures when you deploy. In mission-critical situations, such as defense communications, automotive collision avoidance, and first responder communications, to name a few, the difference between a robust, high performing network could be the difference between life and death for personnel in harm's way. An asset that provides inferior, incomplete, or intermittent sensor data will directly hinder decision makers on the ground from making optimal, timely, life-saving decisions.

The traditional approach of developing and testing in a pristine lab and later considering real-world network conditions after deploying is short-sighted, risky, time consuming, and costly. Thankfully these problems are avoidable.

Aukua's MGA2510 helps mitigate project risk and cost while simultaneously enabling you to develop more robust, superior performing satellite networks. In addition, with the MGA2510 you can resolve post-production issues faster and conduct rapid prototyping under realistic network conditions. Aukua can easily recreate real-world network conditions such as latency, asymmetric bandwidth limits, network congestion, or other negative scenarios.

AUKUA
SYSTEMS

www.aukua.com
+1 866-533-2215