

IMPROVING STARLINK'S LATENCY


Starlink engineering teams have been focused on improving the performance of our network with the goal of delivering a service with stable 20 millisecond (ms) median latency and minimal packet loss.

Over the past month, we have meaningfully reduced median and worst-case latency for users around the world. In the United States alone, we reduced median latency by more than 30%, from 48.5ms to 33ms during hours of peak usage. Worst-case peak hour latency (p99) has dropped by over 60%, from over 150ms to less than 65ms. Outside of the United States, we have also reduced median latency by up to 25% and worst-case latencies by up to 35%.

Latency

Latency refers to the amount of time, usually measured in milliseconds, that it takes for a packet to be sent from your Starlink router to the internet and for the response to be received. This is also known as “round-trip time”, or RTT. Latency is one of the most important factors in your perceived experience when using the internet – web pages load faster, audio and video calls feel closer to real-life, and online gaming feels responsive. As an example, testing has shown that increasing bandwidth beyond around 10 Mbps does not increase web-page load time, but a reduction in latency results in substantially lower load times.

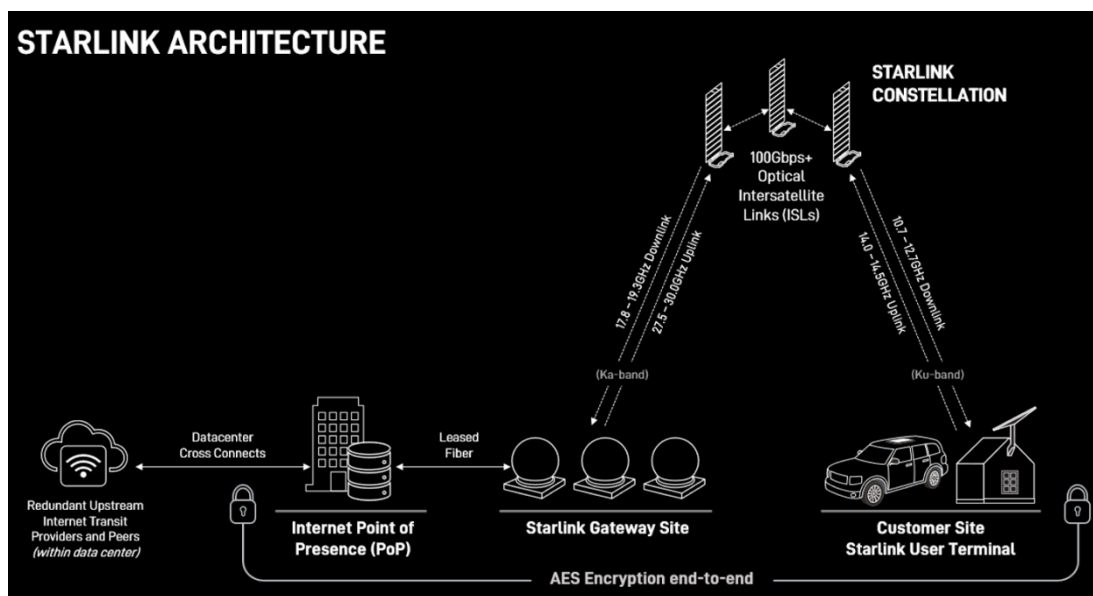
To measure Starlink’s latency, we collect anonymized measurements from millions of Starlink routers every 15 seconds. These 15 second average latencies are then used to calculate the median and worst-case latency. The median (50th percentile or p50) refers to the point where half of the latency measurements are below that number and the other half are above. The worst-case latency, or 99th percentile, is defined as the place where 99% of measurements are better than the point. While we look at data from all points in time, we specifically focus on the performance during hours of peak usage (6-9 PM local time), when the largest number of people are using Starlink, and the network is under the most load.



What Drives Latency

Latency in the Starlink network is driven by several factors. The biggest ones are:

- Physical speed-of-light propagation from the user to the satellite and back to the ground. This is in the range of 1.8-3.6ms per leg, and usually under 10ms for the round-trip. Additional latency can be induced if traffic flows over laser links, instead of directly to the ground (as a result of congestion mitigation, lack of satellite to ground paths, or other factors). While laser connectivity is essential for connecting the most remote locations on Earth and for routing around congestion in the network, we are making strides to ensure that latency sensitive traffic can flow over the shortest path possible.
- Ground latency from the gateway sites to the internet connection point driven by ground network layout. In 2024, we are adding 6 internet connect locations (called Points of Presence, or PoPs) in the US and are optimizing gateway locations and our planning algorithms to ensure that traffic can land as close to its destination point as possible. We will continue to ensure that users are allocated to optimal internet connection locations, so that all users get the lowest latency possible route to the internet.
- Fronthaul (the radio links between the satellite and user) scheduling latency driven by the network topology and the number of users served by a given beam from a satellite. While this latency is an inherent part of shared wireless systems, there is significant room for optimization, and this has been a major focus in the past several months.
- Dumb stuff driven by non-physical limitations in our system – unneeded processing delays, unoptimized buffers, or unnecessary packet drops that force retries. Buffers across our network have been right sized to reduce bufferbloat, and queueing algorithms have been improved to increase capacity on our gateway links from the ground to satellites. Our WiFi latency has been improved, with the addition of active queue management, fq_codel, to the Starlink WiFi router. With active queue management enabled, when one person on your WiFi is downloading a big file, and another is playing a game, the game latency will not be affected by the download.





Continuing Towards Our 20ms Goal

Over the past several months, monitoring and metrics have also been added across the network to measure latency on every subsystem down to the microsecond. We have rigorously tuned our algorithms to prefer paths with lower latency, no matter how small the difference and to remove any and all sources of unnecessary and non-physical latency.

This is just a selection of some of the most impactful changes we have made. Since the beginning of the year, teams have deployed and tested 193 different satellite software builds, 75 gateway software builds, 222 Starlink software builds, and 57 WiFi software builds.

Thank you to our 2.6M+ customers for choosing Starlink. You can expect latency to continue to improve over the coming weeks and months as we prioritize software changes, build additional ground infrastructure, and launch more satellites. In future updates, we'll communicate other performance statistics and goals of our network as we work to improve your experience.

Be sure to check the latest latency statistics for your region at starlink.com/map.