CLOUDBRINK

Packet Loss Uncovered:

The CxO Guide to Boost Remote Network Performance

Packet Loss Uncovered: The CxO Guide to Boost Remote Network Performance

The emerging trends of a growing hybrid workforce and an increase in connected devices at home and remote sites are putting a strain on the underlying network resources.

Businesses need to support their staff, partners, and customers who face challenges with slow and unreliable internet connections. These issues not only strain company resources but also cause widespread frustration.

Everyone expects smooth video communication as a standard of doing business. This may not be possible all the time, especially on wireless or in remote or congested areas.

While latency limits the ability to exploit available bandwidth, packet loss is the dominating factor that severely limits throughput. Even in places where the mid and first mile are typically well provisioned, last mile packet loss continues to be the major factor affecting throughput.

This deep-dive guide will explain:

- The mechanics and implications of packet loss
- How packet loss can affect throughput
- Additional latency contributors that increase connectivity burden
- Mitigation and solutions for packet loss and latency

The Mechanics of Packet Loss

What is packet loss? It is the failure of datagrams to reach their intended destination after transmission across a network. A packet loss below 0.5% (one-half of one percent) is generally recommended for optimal video performance. Packet loss has a significant impact on real-time communication applications, particularly video, but packet loss is different for wired vs wireless networks.

- **Wired Networks:** Packet loss is commonly attributed to various network congestion challenges, software bugs, faulty network hardware, over-subscription, overloaded devices, and various other factors that can result in dropped packets during data transmission.
- **Wireless Networks:** Additional challenges such as multipath interference, co-channel interference, adjacent channel interference, low signal-to-noise ratio, near/far issues, hidden node problems, channel overload, and several other factors are typical. Consequently, wireless networks tend to experience a higher rate of packet loss compared to wired networks.

There are other activities when low network throughput will impact business operations. Activities such as large file transfers that should take minutes could take hours.Efficiently downloading content, accessing file shares, and processing actions in data-intensive applications are critical to maintaining productivity, meeting deadlines, and ultimately, impact the business's financial performance.

When packets are lost, network or application protocols often initiate retransmission requests. For example, in a campus network or in-office scenario where the client is close to the server (typically within 2 msec), these retransmissions are usually fast and imperceptible in most applications.

However, as the distance between the user and the applications increases, retransmissions can cause issues in applications such as VoIP leading to unusable communication due to out-of-order packets and delays in arrival. This is more significant in transport protocols such as TCP, when the distance between user and application increases. This is usually measured in terms of Round Trip Time (RTT). TCP clamps down aggressively on throughput (window size is decreased) and upon recovery of loss conditions open up the window and hence increase throughput slowly.

This can result in significant underutilization of bandwidth. Such protocols end up punishing lossy conditions more punitively because the assumption is that all congestion is due to oversubscription and overflow of queues in network devices. The clamping down of throughput is to relieve the overflowing queues from such congestion.

Another aspect of long latency resulting in lower throughput is due to the limitation of buffer queue sizes in the sender and receivers. The send and receive queues need to be able to hold at least a window worth of data which is a function of bandwidth and delay product. So, latency places an upper bound on the achievable throughput. Additionally, packet loss on top of high latency links results in significantly reduced throughput.

This simplified diagram explains how modern hybrid workers—whether occasionally or always working remotely—often encounter significant packet loss issues on the Wi-Fi networks they utilize at home, hotels, or other locations.

Packet Loss for Hybrid Workers on Wi-Fi

Access
- Home
- Coffee Shops
- Hotels
- Airports

Last Mile
- Broadband Fiber
- DSL
- Cable
- 4G/5G

Mid Mile
- Public Internet
- Peering Agreements

First Mile
- Private Datacenter
- SaaS Provider

Highest Packet loss → **Lower Packet Loss**

Packet loss due to Oversubscription

Network congestion can be experienced when internet broadband connections see an uptick in packet loss due to oversubscription. This occurs when an Internet Service Provider (ISP) sells the same bandwidth to multiple users under the assumption that not all clients will utilize the full bandwidth simultaneously.

For instance, a 1Gb service offered by an ISP might be sold to hundreds of other users, but the network connecting all these hundreds of users to the internet may only be a 10Gb network. Consequently, when heavy usage happens simultaneously, it may cause sporadic spikes in network activity, which in turn contributes to packet loss.

The level of over-subscription often varies based on the provider and the type of service. In general, residential customers experience higher oversubscription ratios, often at 1:100 or more, while for business users, the ratio tends to be lower, around 50:1 or even less.

Packet loss due to outdated mitigation protocols

To mitigate congestion, the internet employs intricate routing protocols. However, when congestion is unavoidable, the network protocols discourage retries or the sending of duplicate packets. Instead, they request that the end user reduce demand and slow down, known as a 'back-off'. The core protocols were designed before the advent of wireless networking when the internet was much smaller.

While this strategy effectively manages core internet overload, it can cause unnecessary connection slowdowns when the issues are more local on the home network or broadband connection. These protocols which work end-to-end can also be overly sensitive to temporary disruptions.

For instance, if your neighbor uses a wireless garage door opener and causes a momentary glitch in your Wi-Fi, resulting in some lost packets, the protocols may initiate a back-off. This occurs despite there being sufficient bandwidth, and the glitch being nothing to do with available network bandwidth.

You can read more on these topics at Remote Network Performance - A Hard Nut to Crack.

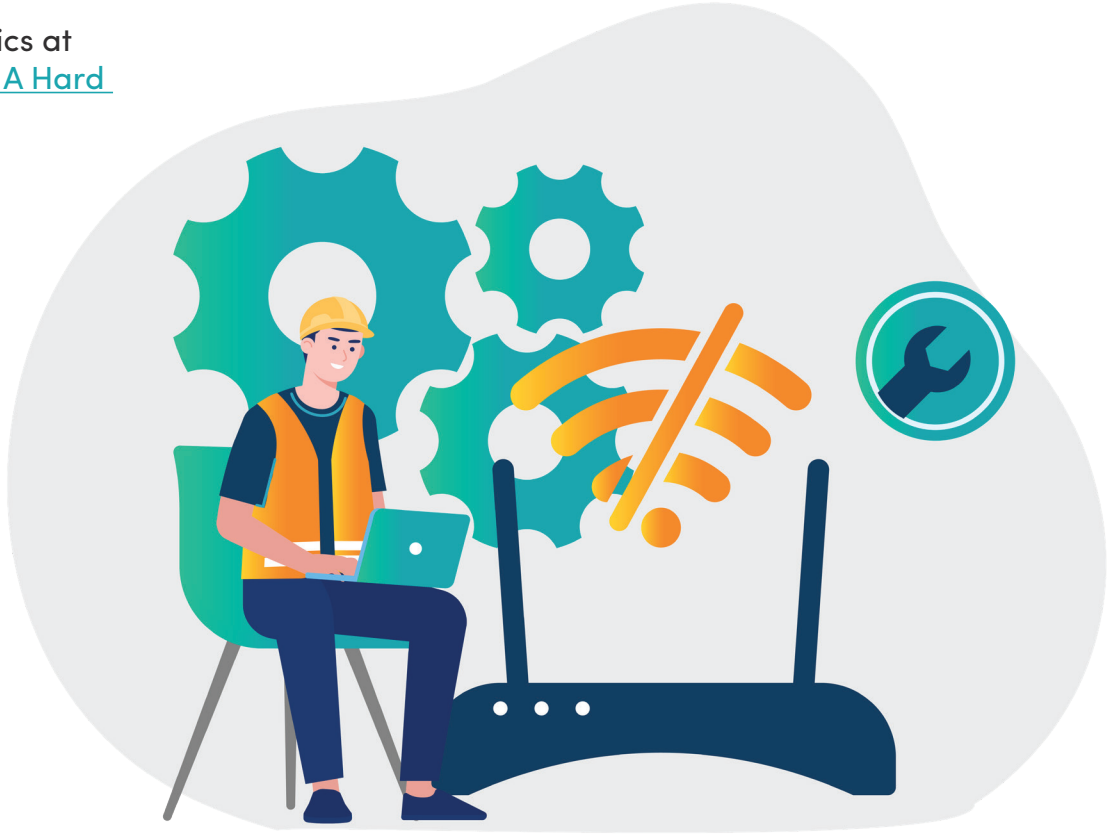

Throughput, Latency, and Packet Loss

While there are several factors that influence throughput, bubbling to the top is latency. As the distance between you and the server or application with which you're communicating increases, throughput begins to decline.

The two charts below compare the impacts of latency and packet loss on throughput.

- The first indicates the results of latency only on throughput, assuming no packet loss. Essentially, anytime latency increases, throughput will be reduced.
- The second chart adds packet loss of 0.0046% to the latency metric, illustrating that any volume of packet loss can have a significant impact on throughput.

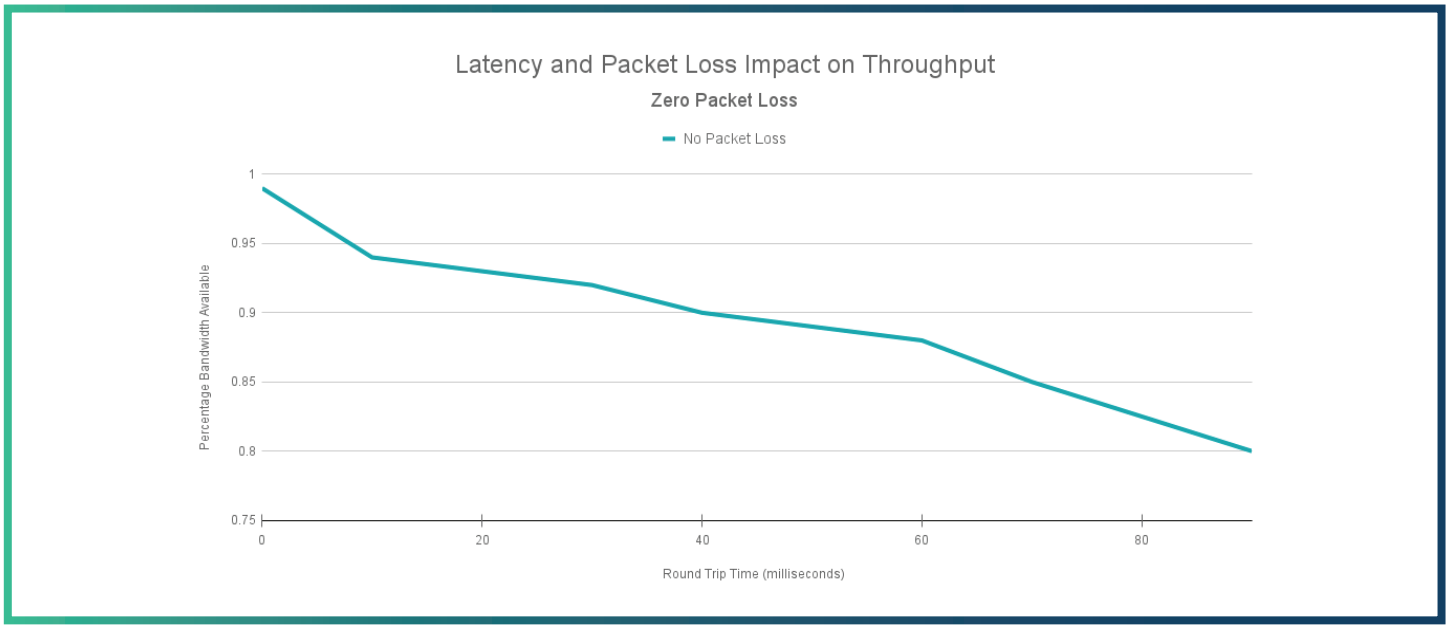

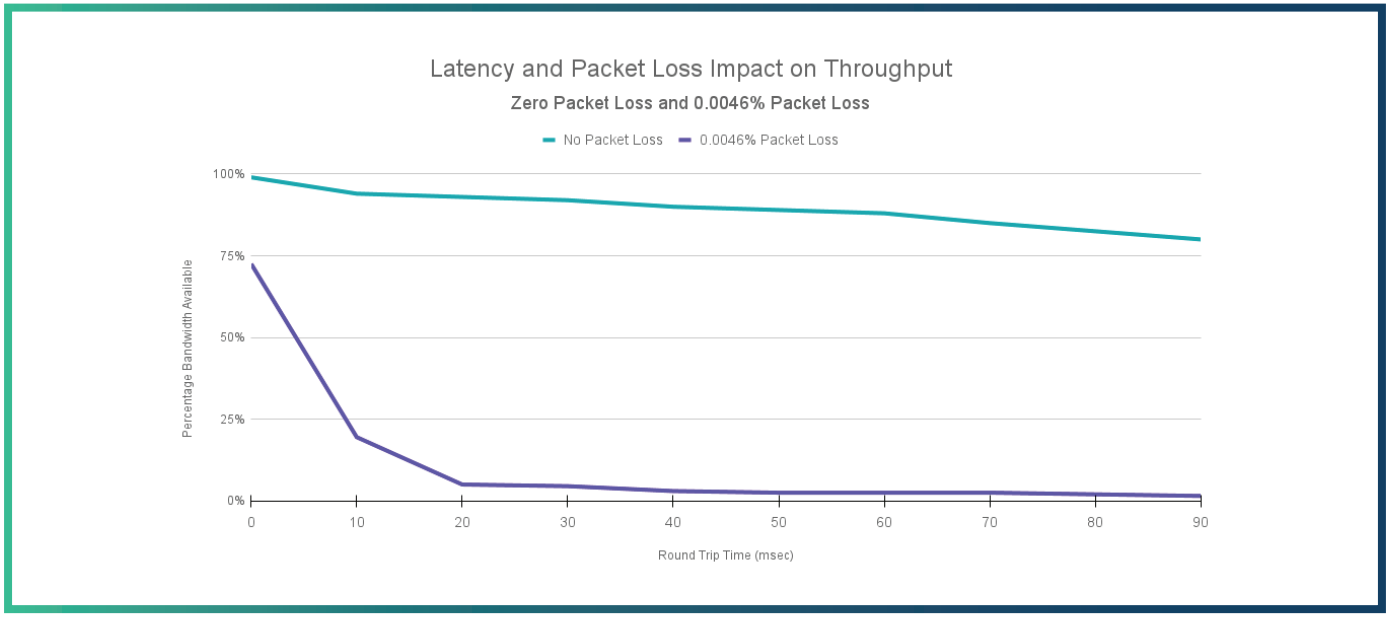

Below are noted some significant data points:

- By adding packet loss to latency, at 60 milliseconds the throughput has notably degraded from an approximately 10% reduction to a 97% reduction.
- With a only one-half percent packet loss and 10 milliseconds latency, throughput can plummet over 90% to leave less than 10% of the available bandwidth.
- When latency surpasses 30 milliseconds, throughput can drop to only 2.5 % of the available bandwidth.

Controlled Environments

In locations like buildings and campus networks designed with care, packet loss on a wired network can be minimized. Likewise, with careful design and the use of appropriate tools for Access Point placement, typical packet loss on Wi-Fi can also be reduced. This suggests that even when services or applications are 50 to 100 milliseconds away, the performance of the network will generally remain satisfactory.

Uncontrolled Environments

In locations like homes, hotels, multi-tenant buildings, airports, and coffee shops, where numerous wireless networks and other wireless devices are competing for the same airwaves, packet loss is common. This, when combined with latency, leads to a severe reduction in usable service.

Top Contributors to Latency

After understanding the effects of latency and its relationship with packet loss, it's important to ensure all sources of latency are identified and addressed.

Distance

The most evident cause is tied to the laws of physics. The greater the distance between you and the applications you're accessing, the higher the latency. The chart below illustrates the relative simplicity of this.

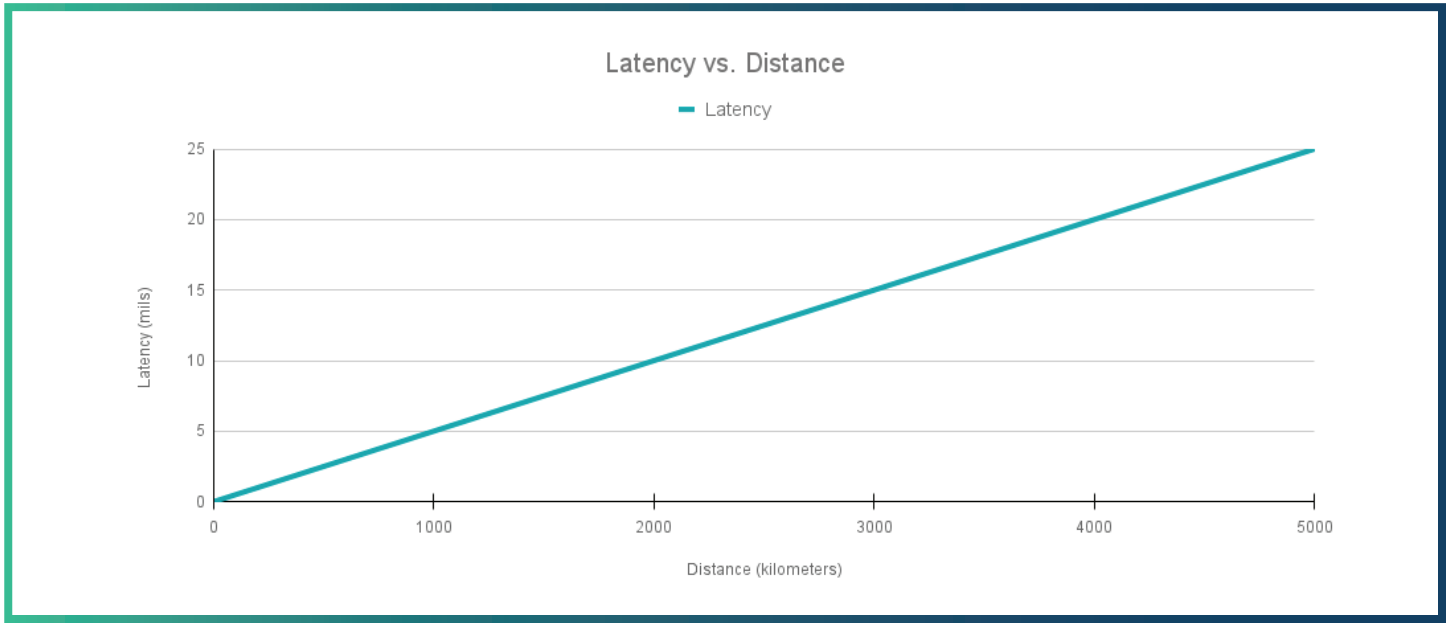

Data Path

A significant latency source arises from the path your data traffic follows. For example, when using a VPN, some or all your data traffic must be routed through the VPN server, often housed at an organization's headquarters. This introduces a few challenges.

- The VPN has a finite capacity, so if an excessive number of users connect concurrently there can be a delay in processing packets.
- This configuration can result in 'hairpinning' for remote users, where data unnecessarily travels extensive distances. This substantially increases latency.

Hairpinning example: A user visiting San Francisco accessing a SaaS app in San Francisco via a VPN server in their New York HQ, causing their data to traverse the entire country and back. This increases latency from under 10 milliseconds to over 100 milliseconds.

Service Provider Peerings

These are agreements where ISPs permit other ISPs to utilize their network, typically routing traffic based on cost-efficiency rather than the fastest route. This practice can heighten latency.

In-line security devices

These can be one of the most significant causes of latency. You might have transitioned to a secure services edge (SSE) solution with a Zero Trust Network Access (ZTNA) to alleviate VPN latency problems. However, this move can also introduce additional latency.

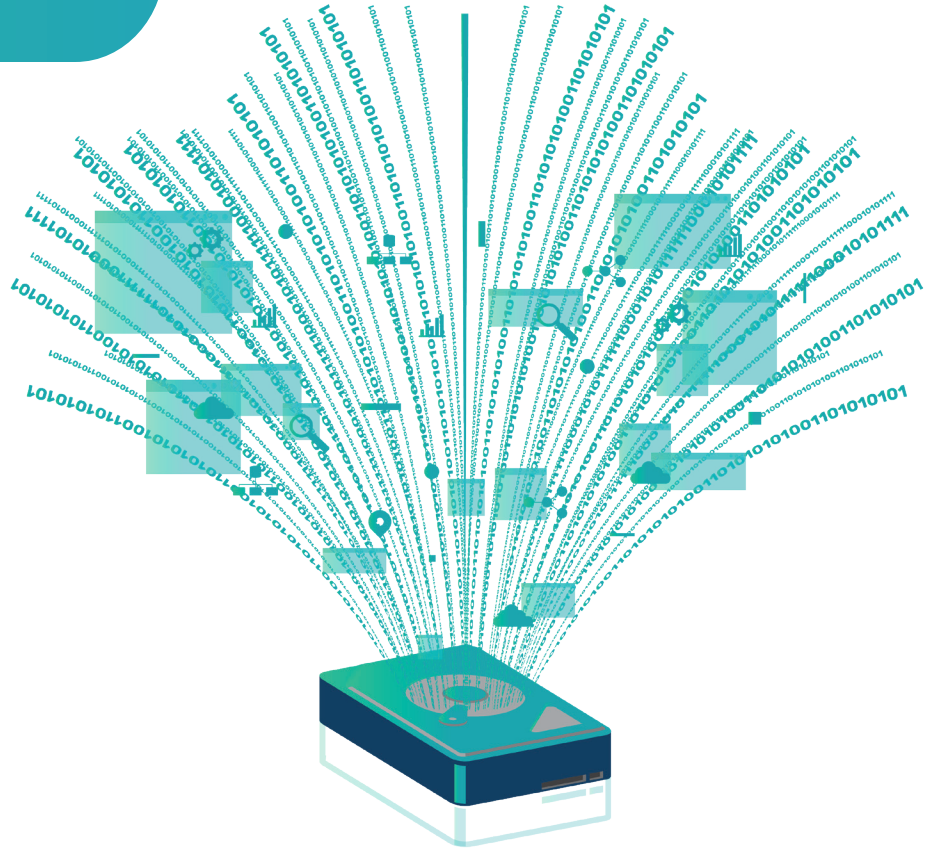

For instance, the service level agreement might stipulate an average latency over a calendar month of 100 milliseconds or less for the 95th percentile of the traffic.

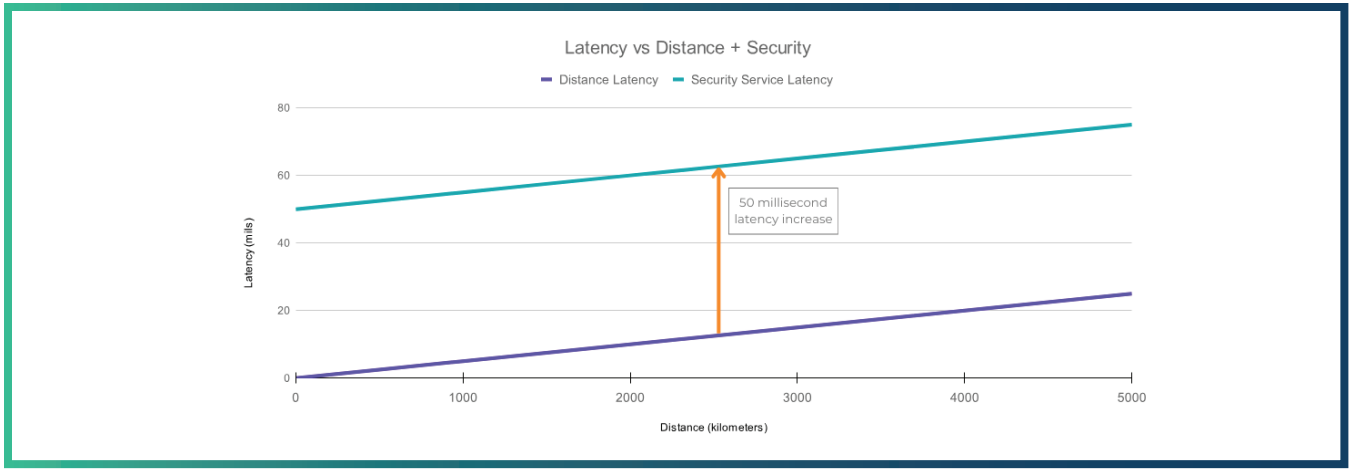

The next chart shows this in more detail by indicating the 50ms starting point with and without packet loss. As you can see, packet loss can have an immediate impact on usable bandwidth.

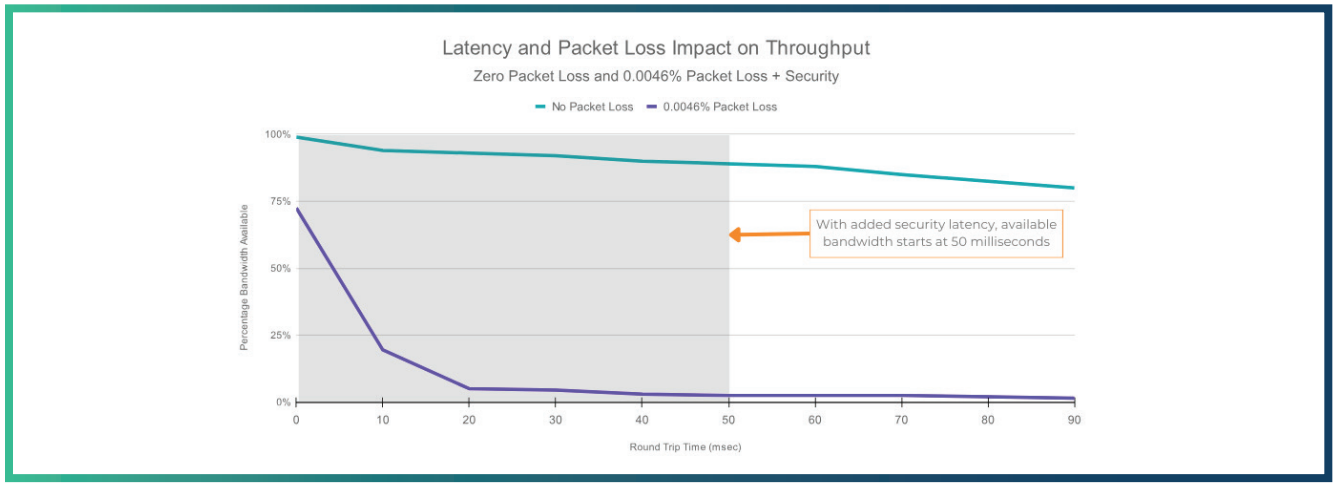

Considering that significant service degradation occurs with just .005% packet loss combined with latency exceeding 10 milliseconds, these types of services are almost certain to cause severe throughput disruption for anyone in a hybrid work environment.

With a conservative delay of just 50 msec - half of the SLA of a major ZTNA vendor - the latency in a local metro region is increased by over 20 times. This immediately puts the end user in the situation where very low packet loss will reduce the performance of the network by over 95%.

How to to Overcome Packet Drops and Latency

It's a known issue that innovation moves fast and often technology doesn't keep up. The perfect storm of increasing remote workers and applications rapidly moving to the cloud is showing a marked increase in packet drop and latency, which in turn leads to poor network and application performance. Legacy network protocols and layered security services were not designed for the modern cloud-centric hybrid work scenario.

Cloudbrink analysis

The team at Cloudbrink has extensively analyzed this problem and designed a groundbreaking technology stack from the ground up. Cloudbrink's service addresses latency and packet loss challenges that can impact network throughput and application responsiveness by using FAST edges and a strategy of pre-emptive and accelerated packet recovery.

FAST edges are software instances positioned close to users. They reduce the hop distance data needs to travel, significantly diminishing end-to-end latency. These edges, coupled with the self-reliant network driven by the Cloudbrink Protocol, enable per-hop recovery, minimizing the distance impaired traffic traverses through the network.

The Cloudbrink Protocol, along with artificial intelligence and machine learning, plays a key role in pre-emptive and accelerated packet recovery. This protocol carries context and makes per-hop decisions to optimize performance and recover lost packets without disrupting end-to-end encryption.

Integrated security and network protocols are a requirement to overcome the impact of latency and packet loss. When distinct security protocols run on top of end-to-end network protocols, they can introduce additional latency and lead to packet drops, affecting network throughput and application responsiveness.

This impact on network throughput and application responsiveness has a domino effect, leading to application delays and slow performance. Such issues can degrade productivity and employee work satisfaction, underscoring the importance of an integrated solution like Cloudbrink.

What Cloudbrink delivers

Cloudbrink is proven to deliver up to **30 times the performance** when compared to other VPN and ZTNA vendors. **Download the full report**.

It's purpose-built to deliver the industry's highest performance connectivity to remote and hybrid workers, anywhere in the world. The Cloudbrink service powers remote workforce connectivity with a simple implementation that solves latency, packet loss, and security challenges by using FAST edges combined with preemptive and accelerated packet recovery.

With Cloudbrink you'll experience:

- Greater security with a reduced attack surface
- Higher Performance
- Lower Cost
- Easier Management
- Faster Implementation

Ready to learn more?

Video How Cloudbrink Works

Visit the Cloudbrink website

Have questions? Request an appointment

About Cloudbrink: Cloudbrink provides a secure, high-performance cloud networking service that significantly improves remote employee productivity by delivering up to a 30X increase in network performance. The company uses AI and ML to provide edge-native Zero-Trust Access for users and devices. Cloudbrink delivers accelerated performance for cloud, SaaS, and data center applications across multi-cloud networks. Cloudbrink's software-only solution includes the world's first high-performance ZTNA with personal SD-WAN and Automated Moving Target Defense (AMTD) security. With the ability to use thousands of dynamically provisioned PoPs called FAST edges, Cloudbrink provides an in-office experience to all workers. This powerful experience for users comes with reduced operational complexity for network, security, cloud and IT administrators.

Charts based on data from Energy Sciences Network and other sources.


CLOUDBRINK