Optimizing your network

A complete guide to understanding jitter, latency and packet loss

Unified communications and collaboration (UCC) is changing the world and the way we work. The worldwide implementation of VoIP and video as major communication solutions is making these changes possible.

But all new technologies come with challenges and one of the major hurdles that IT teams face is network performance. When network performance is compromised, smooth communication becomes problematic and end users have a poor experience.

This comprehensive guide will explain everything you need to know about three common problems experienced as a result of poor network performance: jitter, latency and packet loss.

We'll take an in-depth look at these issues, the reasons you experience them and how to address these issues by optimizing your network.

Packet loss

What is packet loss?

In any network environment, data is sent and received across the network in small units called packets.

This applies to everything you do on the internet, from emailing, uploading or downloading images or files, browsing, streaming, gaming – to voice and video communication. According to a 2017 survey from Statista, in 2017, 24% of surveyed companies claimed that downtimes cost them between $301,000 and $400,000. In most cases, these situations of downtimes might have arisen from a seemingly simple issue that escalated into significant setbacks.

When one or more of these packets is interrupted in its journey, this is known as packet loss. The Transmission Control Protocol (TCP) divides the file into efficiently sized packets for routing. Each packet is separately numbered and includes the destination's internet address. Each individual packet may travel a different route, and when they have arrived, they are restored to the original file by the TCP at the receiving end.

What causes packet loss?

Network congestion

The primary cause of network packet loss is congestion. All networks have space limitations, so in simple terms, network congestion is very much the same as peak hour traffic.

Think of the queues on the road at certain times of the day, like early mornings and the end of the working day. Too much traffic crowding onto the same road can become bottlenecked when it tries to merge, and the result is that it doesn't reach its destination on time.

At peak times, when network traffic hits its maximum limit, packets are discarded and must wait to be delivered. Fortunately, most software is designed to either automatically retrieve and resend those discarded packets or slow down transfer speed.

Network hardware problems

The speed with which hardware becomes outdated or redundant these days is another major problem for your network. Hardware such as firewalls, routers, and network switches consume a lot of power, and can considerably weaken network signals. Sometimes organizations overlook the need to update hardware during expansions or mergers and this can contribute to packet loss or connectivity outages.

Software bugs

Closely related to faulty hardware is a buggy software running on the network device. Bugs or glitches in your system can sometimes be responsible for disrupting network performance and preventing the delivery of packets. Hardware reboots and patches may fix bugs.

Overtaxed devices

When a network is operating at a higher capacity than it was designed to handle, it weakens and becomes unable to process packets, and drops them. Most devices have built-in buffers to assign packets to holding patterns until they can be sent.

Wired packet loss vs wireless packet loss

As a rule, wireless networks experience more issues with packet loss than wired networks. Radio frequency interference, weaker signals, distance and physical barriers like walls can all cause wireless networks to drop packets.

With wired networks, faulty cables can be the culprit, impeding signal flow through the cable.

Security threats

If you're noticing unusually high rates of packet drop, the problem could be a security breach. Cybercriminals hack into your router and instruct it to drop packets. Another way that hackers can cause packet loss is to execute a denial-of-service attack (DoS), preventing legitimate users from accessing files, emails, or online accounts by flooding the network with too much traffic to handle. Packet loss can be difficult to fix during a full-blown security.

Deficient infrastructure

This highlights the importance of a comprehensive network monitoring solution. Some out-of-packet monitoring tools were not engineered for the job they've been assigned to do and have limited functionality.

The only way to effectively deal with packet loss issues is to deploy a seamless network monitoring and troubleshooting platform that can view your entire system from a single window. In a nutshell, comprehensive network monitoring solution = packet loss fix.

Ping and packet loss

When it comes to the determining what constitutes a strong internet connection, and the reduction of random packet loss, there are three factors to consider: upload speed, download speed and ping.

Upload speed
This is how fast you can send data to others. Uploading is used when sending large files through email, or in using video to chat with others. Upload speed is measured in megabits per second (Mbps).

Download speed
This is how fast you can pull data from the server to you. By default, connections are designed to download more quickly than they upload. Download speed is also measured in Mbps.

Ping
This is the reaction time of your connection, or how quickly you get a response after sending out a request. A fast ping means a more responsive connection, and this is especially important in real-time applications like gaming, and voice and video calls. Ping is measured in milliseconds (ms).

Anything below a ping of 20 ms is considered ideal, while anything over 150 ms would result in noticeable lag.

Even though your ping is good you may still be having issues with packet loss. Although the data is being sent and ultimately received quickly by the destination server, some data might not be getting there correctly.

The effects of packet loss

For users, packet loss can be more than annoying, particularly in real-time processes like VoIP and video conferencing. According to a QoS tutorial by Cisco, packet loss on VoIP traffic should be kept below 1% and between 0.05% and 5% depending on the type of video.

Different applications are affected by packet loss in different ways. For example, when downloading data files, a 10% packet loss might add only one second to a ten second download. If packet loss rate is higher, or there is high latency, it can cause delays to be worse.

Real-time applications like voice and video will be affected more severely by packet loss. Something as small as a 2% packet loss is usually quite noticeable to a listener or viewer, and can cause the conversation to be stilted and unintelligible.

The effects of packet loss also differs depending on the application/protocol (TCP/UDP). If a packet is dropped, or not acknowledged, TCP protocol is designed to retransmit it. UDP, however, doesn't have the capability to retransmit, and therefore doesn't handle packet loss as well.

Diagnosing and fixing packet loss

Everyone has experienced packet loss in voice calls. This is where comprehensive network monitoring and troubleshooting comes into its own. Network monitoring can quickly and accurately diagnose and identify the root causes of packet loss problems such as in the following examples.

Example 1:
During a Teams call, the quality deteriorates and becomes distorted and patchy, or eventually drops out completely. But even though Teams may be having issues, you might still be able to successfully communicate using Cisco Webex, Google Hangouts or WhatsApp. This is because of the difference in the way that each specific program transmits over the internet, and the route that the packets take.

Example 2:
You may be on a call with a perfect connection to a server in Springfield, IL but then find you're experiencing an exceptionally high packet loss when connecting to a server in Richmond, VA. This would indicate problems with the pipeline between your location and the server in Richmond.

Do a ping test
A ping test is a diagnostic tool that provides data on how well an internet-enabled device communicates with another endpoint. A ping test can assess network delays or issues by sending an Internet Control Message Protocol (ICMP) packet – or ping – to a specific destination.

ICMP packets contain only a tiny amount of information, so they don't use much bandwidth. When the ping reaches the device, that device recognizes and replies to the originating device. The total time taken for the ping to arrive and return is recorded as 'ping time' or 'round trip time'.

If the number of packets sent and received are not equal, this means some packets never arrived to or from your phone. This inevitably leads to call quality issues like choppy voices, extended silences, jumbled audio and other call quality problems.

Deep packet inspection
Any organization with a private network will have hundreds or even thousands of unique connections and data transfers every day.

Deep Packet Inspection (DPI) is an in-depth way of examining and managing network traffic. DPI is one of the most important tasks that network administrators need to carry out. It locates, identifies, blocks or re-routes packets with specific data or code. It examines the contents of packets passing through a given point and determines what the packet contains. Most network packets are split into three parts:

Header – containing instructions about the data carried by the packet such as length, synchronization, packet number, protocol as well as originating and destination addresses.

Payload – the actual data contents, or body of the packet.

Trailer – also referred to as the footer tells the receiving device that it has reached the end of the packet.

Traceroute packet loss and high latency
Traceroute is a command-line tool that comes with Windows and other operating systems. Along with the ping command, it's an important tool for understanding Internet connection problems.

If you're having trouble connecting to a website, traceroute can tell you where the problem is. It can also help visualize the path traffic takes between your computer and a web server.

Monitoring packet loss

Every network experiences some degree of packet loss, but what is acceptable? The most important thing to remember is that prevention is better than cure when implementing packet loss solutions.

Network monitoring should be the first strategy you use to preserve and uphold the integrity of your network environment. Regularly scanning your devices will ensure that your routers are capable of handling capacity, and your system is equipped to prevent data loss.

Latency

What is network latency?

Network latency, sometimes called lag, is the term used to describe delays in communication over a network. In networking, latency is best thought of as the amount of time it takes for a packet of data to be captured, transmitted, processed through multiple devices, then received at its destination and decoded.

When delays in transmission are small, it's referred to as a low-latency network (desirable) and longer delays are called a high-latency network (not so desirable).

Long delays that occur in high-latency networks create bottlenecks in communication. In the worst cases, it's like traffic on a four-lane highway trying to merge into a single lane. High latency decreases communication bandwidth, and can be temporary or permanent, depending on the source of the delays.

Latency is measured in milliseconds, or during speed tests, it's referred to as a ping rate. Obviously, zero to low latency in communication is what we all want. However, standard latency for a network is explained slightly differently in various contexts, and latency issues also vary from one network to another.

Causes of network latency

Distance

One of the main causes of network latency is distance, or how far away the device making requests is located from the servers responding to those requests.

For example, network latency between cities: if a website is hosted in a data center in Trenton, New Jersey, it will respond faster to requests from users in Farmingdale, New York (100 miles away), or most likely within 10-15 milliseconds. On the other hand, users in Denver, Colorado (about 1,800 miles away) will face longer delays of up to 50 milliseconds.

The amount of time it takes for a request to reach a client device is referred to as Round Trip Time (RTT). While an increase of a few milliseconds might seem negligible, there are other considerations that can increase latency, such as:

- The to-and-fro communication necessary for the client and server to make that connection in the first place.
- The total size and load time of the page.
- Problems with network hardware which the data passes through along the way.

Data travelling back and forth across the internet often has to cross multiple Internet Exchange Points (IXPs), where routers process and route the data packets, often having to break them up into smaller packets. All this additional activity adds a few milliseconds to RTT.

Website construction

The way webpages are constructed makes a difference to latency. Webpages that carry heavy content, large images, or load content from several third- party websites may perform more slowly, as browsers need to download larger files to display them.

End-user issues

Network problems might appear to be responsible for latency, but sometimes RTT latency is the result of the end-user device being low on memory or CPU cycles to respond in a reasonable timeframe.

Physical issues

In a physical context, common network latency causes are the components that move data from one point to the next like physical cabling such as routers, switches and WiFi access points. In addition, latency can be influenced by other network devices like application load balancers, security devices, firewalls and Intrusion Prevention Systems (IPS).

Other types of latency

The following describes two other examples of the effects of latency.

Fiber optic latency
In the case of fiber optic networks, latency refers to the time delay that affects light as it travels through the fiber optic network. Latency increases over the distance traveled, so this must also be factored in to compute the latency for any fiber optic route.

Based on the speed of light (299,792,458 meters/second), there is a latency of 3.33 microseconds (0.000001 of a second) for every kilometer covered. Light travels slower in a cable which means the latency of light traveling in a fibre optic cable is around **4.9 microseconds per kilometer.**

The quality of fiber optic cable is an important factor in reducing latency in a network.

VoIP latency
The reasons behind audio latency are based on the speed of sound. Latency in VoIP is the difference in time between when a voice packet is transmitted and the moment it reaches its destination. A latency of 20 ms is normal for VoIP calls; a latency of up to 150 ms is barely noticeable and therefore acceptable. Any higher than that, however, and quality starts to diminish. At 300 ms or higher, it becomes completely unacceptable.

High latency in VoIP can severely affect call quality, resulting in:

- Slow and interrupted phone conversations.
- Overlapping noises, with one speaker interrupting the other.
- Echo.
- Disturbed synchronization between voice and other data types, especially during video conferencing.

Reasons behind VoIP latency and how to address them:

Insufficient bandwidth – with a slow internet connection, insufficient bandwidth means that data packets take more time reach their destination, and often arrive in the wrong order.

Firewall blocking traffic – to prevent bottlenecks, always allow clearance for your VoIP applications within your firewall software.

Wrong codecs – codecs encode voice signals into digital data ready to be transmitted. This is often an issue that your provider needs to solve, however some VoIP apps allow you to tweak codecs.

Outdated hardware – sometimes the mix of old hardware and new software can cause latency problems. Changing your telephone adaptor or other VoIP-specific software can help. Even your headset can cause latency.

Signal conversion – if your system is converting your signal to or from analog and digital, this could cause latency.

Best practices for monitoring and improving network latency

The slowing of your network can be extremely problematic in the business world, where time is such a precious commodity. As your network grows bigger, having additional connections means more points where delays and issues can happen.

Problems can increase again as more and more organizations connect to cloud servers, use more applications and expand to accommodate remote workers extra branch offices.

Everyone has experienced latency in various aspects of daily business, and it can severely threaten deadlines, expected outcomes and eventually ROI. This is where comprehensive network monitoring and troubleshooting comes into its own. Network monitoring and troubleshooting can quickly and accurately diagnose and identify the root causes of latency and put solutions in place to reduce impact and improve the problem.

Before you can do anything to improve your network latency, you need to know how to calculate and measure it. By becoming familiar with your latency, you're far better equipped to troubleshoot.

How to check network latency
If you feel that your network is running slow, you can check your latency manually by using Windows.
Open a command prompt and type tracert followed by the destination you'd like to query, such as **cloud.google.com.**

How to measure network latency
Network monitoring and management tools will get this information automatically, but here's how to do it manually. Once you type in the **tracert command**, you'll see a list of all routers on the path to that website address, followed by a **time measurement in milliseconds** (ms). Add up all the measurements, and the **resulting quantity is the latency between your machine and the website in question.**

Latency can either be measured as the Round Trip Time (RTT) or the Time to First Byte (TTFB):

- **RTT** - the amount of time it takes a packet to get from the client to the server and back.
- **TTFB** - the amount of time it takes for the server to receive the first byte of data when the client sends a request.

How to reduce network latency
One simple way to improve network latency is to check that others on your network aren't unnecessarily using up your bandwidth, or increasing your latency with excessive downloads or streaming. Then, check application performance to determine whether applications are acting unexpectedly and potentially placing pressure on the network.

Subnetting is another way to help reduce latency across your network, by grouping together endpoints that communicate most frequently with each other.

Additionally, you could use traffic shaping and bandwidth allocation to improve latency for the business-critical parts of your network.

Finally, you can use a load balancer to help offload traffic to parts of the network with the capacity to handle some additional activity.

How to troubleshoot network latency issues
Manually troubleshooting issues across a large network can become complex, which highlights again the importance of network monitoring and troubleshooting tools.

To check if any of the devices on your network are specifically causing issues, you can try disconnecting computers or network devices and restarting all the hardware. You'll need to ensure that you have network monitoring deployed.

If you still have latency problems after checking all your local devices, it's the issues could be coming from the destination you're trying to connect to.

How to test network latency
Testing network latency can be done by using ping or traceroute (tracert), although, comprehensive network monitoring and performance managers can test and check latency more accurately.

Maintaining a reliable network is an important part of a smoothly operating business. Network issues can become worse if they're not managed properly.

What tools help improve network latency?

Network monitoring and troubleshooting tools are the best way to keep tabs on latency, as well as the other most troubling network problems, packet loss and jitter. You can typically **set network standard expectations for latency and create alerts** when the network latency reaches a certain threshold above this baseline.

Network monitoring tools can help you set up data comparisons between different metrics. This can help you identify performance issues, such as application performance or **errors also affecting network latency.**

A network mapping tool can also help you pinpoint where within the network the performance issues are occurring, which allows you to troubleshoot problems more quickly.

Specific **traceroute tools** monitor packets and how they move across an IP network, including how many "hops" the packet took, the roundtrip time, best time (in milliseconds), as well as the IP addresses and countries the packet traveled through.

By improving your network speed and reducing latency, your business processes will also make leaps and bounds towards efficiency and high performance.

Jitter

What is network jitter?

Technically, jitter is a variance in latency, or the time delay between when a signal is transmitted and when it is received. This variance is measured in milliseconds (ms) and is described as the disruption in the normal sequence of sending data packets.

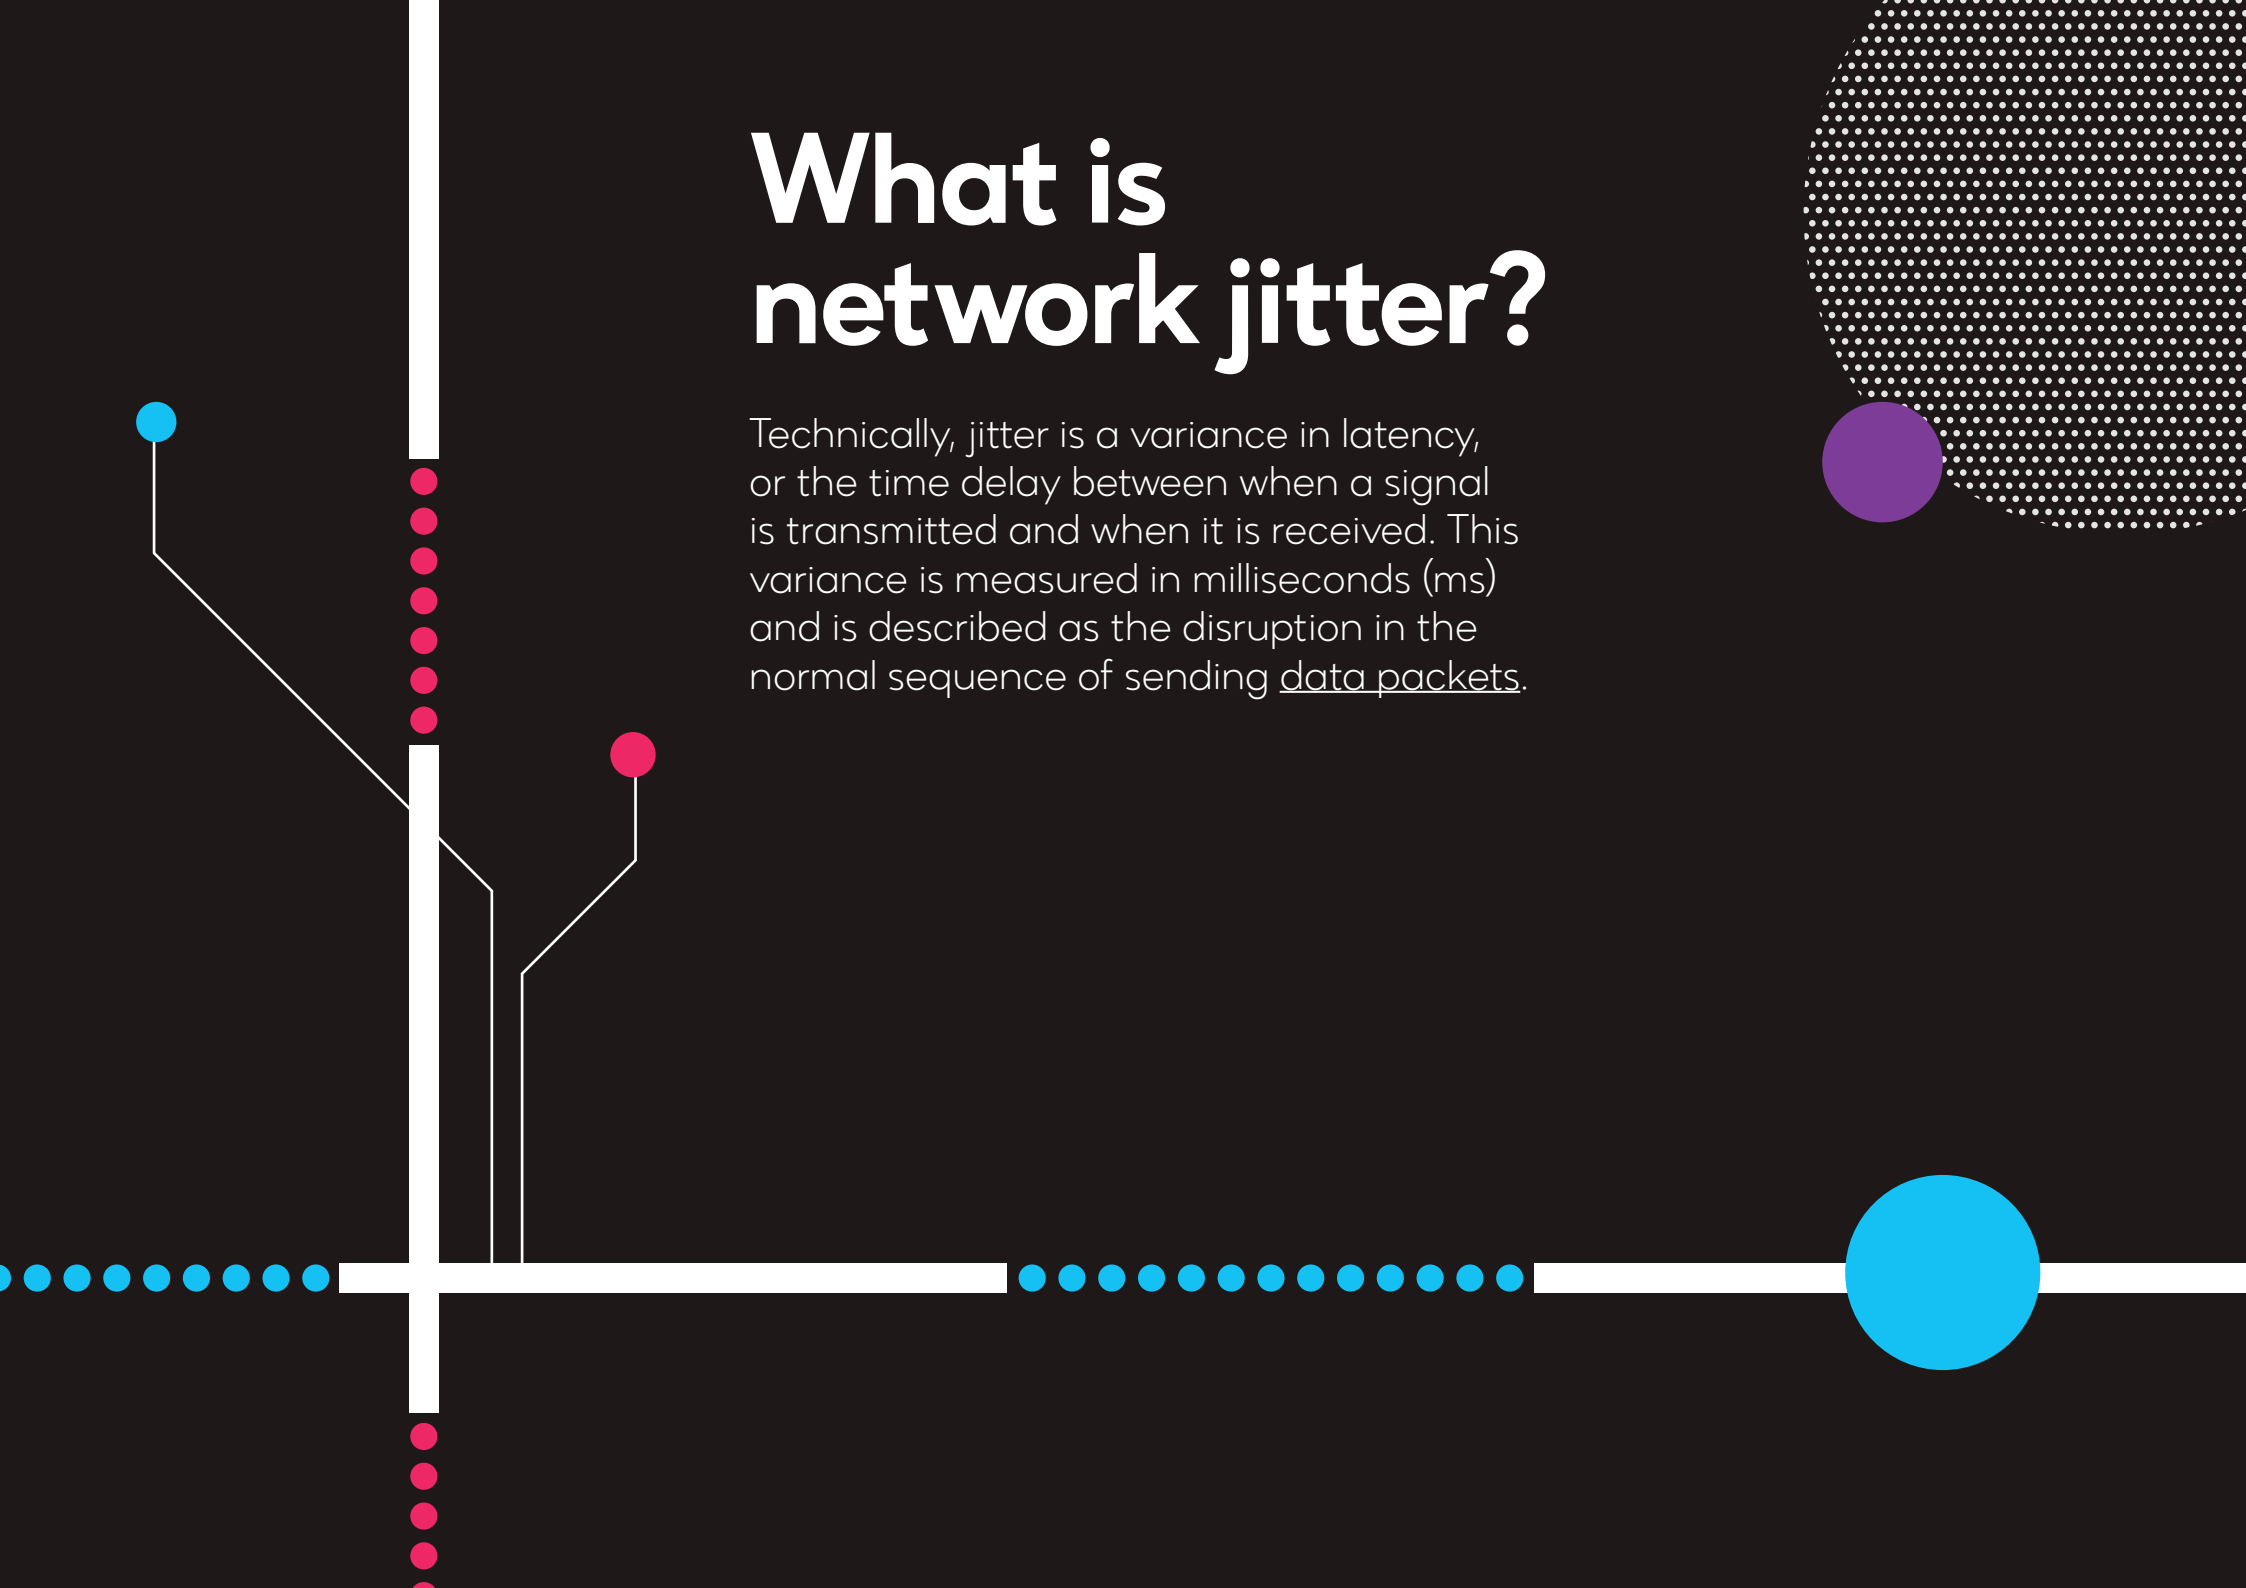

The technology behind jitter

Data packets

To define jitter in networking, it comes down to packets.

Packet jitter definition means that all data, in fact everything you do on the internet, involves packets. All text, images, audio or video is transmitted in packets over a given network path. When you send or receive an email, search for information on web pages, stream, game or shop online, digital information is despatched, received, 'unscrambled' and 'reassembled' ready to view and listen to. Packet switched networks allow the exchange of all this information.

Most network packets are split into three parts:

Header
The header contains instructions about the data carried by packet such as:

- Packet length - some networks have fixed-length packets, while others rely on the header to contain this information.
- Synchronization - to help the packet align with the network.
- Packet number - identifying which packet in the sequence.
- Protocol - defines what type of packet is being transmitted whether it's e-mail, web page, streaming video.
- Destination address - where the packet is going.
- Originating address - where the packet came from.

Payload
This is the actual data or body that is being delivering to the destination. If a packet is fixed-length, then the payload may be topped up with blank information to make it the right size.

Trailer
The trailer, sometimes called the footer, tells the receiving device that it has reached the end of the packet. It may also have some type of error checking. The most common error checking used in packets is Cyclic Redundancy Check (CRC).

Jitter VoIP

VoIP technology converts fragments of your voice into data packets which are transmitted digitally via the internet. One of the most common causes of jitter on VoIP services is the **absence of packet prioritization**. If voice packets aren't prioritized, then the end user is very likely to experience jitter.

A massive amount of information is constantly being transmitted back and forth - millions of packets every second - and all this data takes a toll on network resources, often resulting in delay. The delay may not be as apparent when downloading a file or an email, but when your voice arrives in disorganized packets, it will sound distorted and out of sequence.

Examples of jitter

Constant jitter
This is a roughly constant level of packet to packet delay variation.

Transient jitter
Characterized by a substantial gradual delay that may be incurred by a single packet.

Short term delay variation
An increase in delay that persists for some number of packets and may be accompanied by an increase in packet to packet delay variation. This type of jitter is usually due to congestion and route changes.

The effects of jitter

Packet jitter can cause flickering display monitors, delayed data transmission and poor processor performance.

IP jitter in VoIP communication can severely impact the call quality of telephony and video conferencing, even causing conversations to 'drop out', and become jumbled and difficult to understand.

As a result, high jitter is a big problem for real-time applications like digital voice and video communication, as well as streaming and online gaming.

What is acceptable network jitter?

Some applications and services have a higher level of tolerance for jitter than others. So what is acceptable network jitter?

For example, jitter doesn't affect sending emails as much as it would a voice chat. So, it depends on what we're willing to accept as irregularities and fluctuations in data transfers. But poor audio and video quality leads to a poor user experience and can impact an organization's bottom line.

All networks experience some amount of latency, especially wide area networks. Ideally, over a normally functioning network, packets travel in equal intervals, with a 10ms delay between packets. With high jitter, this could increase to 50ms, severely disrupting the intervals and making it difficult for the receiving computer to process the data.

Ideally, jitter should be below 30ms. Packet loss should be no more than 1%, and network latency shouldn't exceed 150 ms one-way (300 ms return).

What can cause network jitter?

Managing network jitter comes down to understanding what causes jitter in computer networks. Doing a regular network jitter test can reduce the prevalence of jitter within your network.

Network congestion
Insufficient bandwidth is a common problem. Networks become overcrowded with traffic congestion when too many active devices are consuming bandwidth.

Poor hardware performance
Older networks with outdated equipment including routers, cables or switches could be the causes of jitter.

Wireless jitter
One of the downsides of using a wireless network is a lower-quality network connection. Wired connections will help to ensure that voice and video call systems deliver a higher quality user experience.

Not implementing packet prioritization
For VoIP systems in particular, jitter occurs when audio data is not prioritized to be delivered before other types of traffic.

Quality of Service (QoS) and jitter

QoS is the technology that manages data traffic in order to reduce jitter on your network and prevent or reduce the degradation of quality. QoS controls and manages network resources by setting priorities by which data is sent on the network.

There are tools and techniques which are often included in an organization's network Service Level Agreement (SLA) to guarantee an acceptable level of performance.

QoS tools to address jitter

Queuing
Enables you to prioritize or order packets so that delay-sensitive packets leave their queues more quickly than delay-insensitive packets.

Link fragmentation and interleaving (LFI)
Routers do not pre-empt a packet that is currently being transmitted, so LFI reduces the sizes of larger packets into smaller fragments before sending them.

Compression
Payload or headers can be compressed, and this reduces the overall number of bits required to transmit the data. This requires less bandwidth, meaning queues shrink, which in turn reduces delay.

Traffic shaping
Artificially increases delay to reduce drops inside a Frame Relay or ATM network.

How is jitter measured?

Single endpoint
Where your network has control over just one of the endpoints (aka single-ended), jitter is determined by measuring the mean round-trip time (RTT), and the minimum RTT of a series of voice packets.

Double endpoint
In a double-ended path, the measurement used is the instantaneous jitter, or the variation between the intervals for transmitting and receiving a single packet. Jitter is the average difference between instantaneously measured jitter and the average instantaneous jitter throughout the transmission of a series of data packets.

Bandwidth testing
Performing a bandwidth test can also determine the level of jitter. A bandwidth test assesses your internet connection's upload and download speeds, jitter times and your network's overall capacity.

How to fix network jitter issues

Troubleshooting network jitter can be tricky because of its unpredictability. Keeping jitter to a minimum begins by ensuring that your network is initially properly set up. Ensuring a quality network connection, enough bandwidth, and predictable latency can help reduce network jitter.

Jitter buffering
VoIP endpoints such as desk phones and ATAs usually include a jitter buffer to intentionally delay incoming data packets. A jitter buffer ensures that the receiving device can store a set number of packets and then realign them into the proper order, so that the receiver experiences minimum sound distortion.

Jitter buffers are one way to address network jitter and latency but will not always work. If a jitter buffer is too small then too many packets may be discarded, meaning bad call quality. If a jitter buffer is too large, then the additional delay can lead to conversational difficulty.

A typical jitter buffer configuration is 30ms to 50ms in size. You can increase buffer size to a point, but usually they are only effective for delay variations of less than 100 ms.

Perform a bandwidth test
Bandwidth testing sends files over a network to a specific computer, then measures the time required for the files to download at the destination. This determines a theoretical data speed between the two points, measured in kilobits per second (Kbps) or megabits per second (Mbps).

Bandwidth tests can vary greatly. Factors that affect testing can be internet traffic, noise on data lines, file sizes, and load demand on the server at the time of testing. Bandwidth testing should ideally be carried out several times to determine an average throughput.

Improvements from within
Solving your VoIP network jitter problems may not be as challenging as you think.

Upgrade your ethernet cable - Outdated cables and switches can often cause high jitter issues. The latest cables are capable of transmitting data at 250 MHz, as opposed to 125MHz, potentially solving ethernet jitter.

Check your device frequency - A VoIP phone that operates at a higher frequency than a standard 2.4 GHz could cause interference on your network. Some phones run at frequencies as high as 5.8GHz, which could potentially exacerbate jitter across your network.

Reduce unnecessary bandwidth usage during work hours - Using large amounts of bandwidth for activities not related to work, like network gaming, or streaming video content, can make jitter worse.

Schedule updates outside of business hours – Updating applications and operating systems should be carried out outside work times to free up capacity for more essential communications.

Summary

Delivering a great user experience across your UC environment is no small task. Network jitter, packet loss and latency, are major obstacles standing in the way of clear communication and can universally affect your user experience.

But with careful planning and the right technology partners you can easily manage these challenges and optimize your network to enhance the UC experience for your colleagues and customers.

Who are IR?

The modern world relies on a complex array of technologies to keep turning. IR's aim is to simplify that complexity.

We provide insights, monitoring and support to keep your business-critical systems running as they should.

More than 1,000 organizations in over 60 countries rely on IR's experience management solutions.

Contact us

Australia
Tel: +61 (2) 9966 106

USA
Tel: +1 (303) 390 870

UK
Tel: +44 (0) 1895 817 800

Germany
Tel: +49 (89) 97 007 132

Singapore
Tel: +65 6813 0851

ir.com