TELCOWEB

telcoweb.com.br

EBOOK

Addressing bottlenecks in your corporate network

The first step to diagnosing poor performance

SUMMARY

Addressing bottlenecks in your corporate network

One of the biggest challenges for any network administrator is when a user complains about the network performance. Whether due to the wide range of possible causes or the lack of detail from the user, the network administrator is often forced to weigh objective factors, such as actual failures, against subjective ones, such as the user's perception.

This document aims to present a basic technical analysis of this type of complaint, its main causes, and some troubleshooting tips to identify the root cause and resolve the issue.

As mentioned in the first paragraph, the perception of poor performance can stem from various objective causes, ranging from the user's local network to the application design itself, including the communication protocol between these two ends. For example, consider a user watching a video on Netflix. Like YouTube and many other platforms, Netflix uses a video streaming application, which exhibits predominantly unidirectional traffic behavior.

That is, the server application creates a network channel through which video content packets flow to the codec on the other side. On the client side, a system known as a codec (short for coder/decoder) is responsible for receiving packets, temporarily storing them, and converting them into presentation media. Figure 1 provides a basic overview of video streaming architecture.

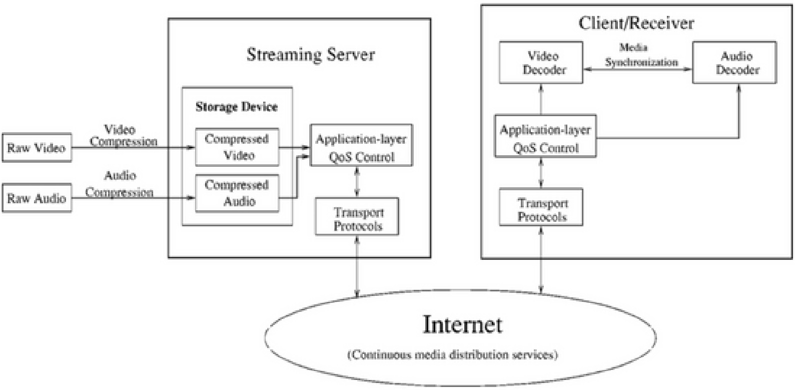

Figure 1 – Video streaming architecture, ref.: DOI:10.5120/909-1287

Latency

In this architecture, which stretches from the streaming source to the end user, the main objective network factor affecting the user experience is latency. High latency—above 200ms—can cause slow buffering rates, interruptions during playback, and delays in video startup.

To reduce latency caused by distance, streaming service providers have deployed network resources such as cache servers. Cache servers are distributed across the infrastructure of Internet providers to bring video content closer to users, thereby reducing latency and improving user experience. In practice, regardless of the video or its provider, it is very likely that the content is being delivered by a cache server.

Packet Loss

Another common complaint about network performance occurs with client-server applications based on databases. In this scenario, the database server is typically located remotely or in the cloud, while the client resides on the local network. The network traffic generated by this type of application is asymmetric-interactive; that is, queries are sent by the client application to the server, which then accesses the database (which may or may not be on the same server).

In addition to network latency, which also significantly impacts system performance, the packet loss rate is a crucial factor. As it increases, user experience deteriorates considerably.

The technical explanation is straightforward: packet loss in the client-to-server direction compromises query integrity, requiring users to resubmit them. In the reverse direction, server-to-client, data loss causes retransmissions which, depending on the data volume, greatly increase the perception of poor performance.

A widely used application that also suffers from packet loss is online video conferencing. The traffic generated by video conferencing is typically symmetrical, since all users are transmitting and receiving audio/video simultaneously. As it is an online event involving relatively high data exchange, packet loss rates above 5% require retransmissions and cause video/audio interruptions or occasional disconnections, which affect all participants. Figure 2 shows the architecture of a typical video conferencing system.

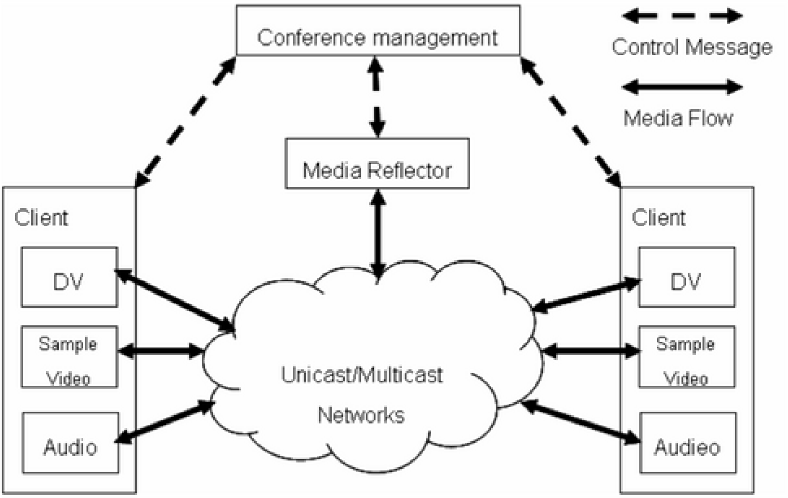

Figure 2 – Video conferencing architecture, ref.: DOI:10.1109/IPC.2007.44

Jitter

Jitter can be understood as the variation in latency and is particularly relevant for Voice over IP (VoIP) communications. Jitter occurs in unstable networks, where packets take varying amounts of time to travel between endpoints. This is usually due to congestion points or routing failures.

Bandwidth

Lastly, bandwidth consumption by modern applications has increased significantly. For fairness, some cloud-based applications and online games simply won't function unless a minimum available bandwidth greater than 20Mbps is provided.

In response to this trend, Internet service providers have increased broadband speeds to meet demand. Plans offering 600Mbps download and 200Mbps upload speeds can easily be found for under R$300/month. However, while increasing bandwidth may immediately improve the user experience, doing so indiscriminately without managing or monitoring system performance can lead to an expensive and ineffective cycle. For example, enabling cloud backup policies during business hours is strongly discouraged.

Methodology for assessing performance complaints

The first step when receiving a slowness complaint is to try to *understand* the user's report more thoroughly. Basic questions to the user can provide clues about possible causes and simplify the analysis process. However, before contacting the user, it's worth taking a quick look at the network performance dashboard, if available. This preliminary check allows for more focused questions, whose answers may confirm an initial diagnosis and provide the user with a clearer position.

When contacting the complaining user, some standard questions can help the administrator understand the scope of the issue. For example, asking whether the slowness occurs only with a specific application or with all applications already provides an important clue about where the failure might be located.

- If the user responds that the issue is with a specific application, it is more likely that the problem lies in the server infrastructure for that application or in its database. If the user reports a general poor performance, further analysis is necessary.
- If the poor performance problem is general, it should be verified whether it truly affects only the complaining user or if it is a broader issue affecting multiple users. If limited to the complaining user, the issue is likely with the PC or its connection to the network. If it's a broader issue affecting multiple users, a deeper investigation is required. At this point, it is crucial that the network administrator has a good grasp of the full range of possibilities offered by the network topology and traffic characteristics. A good way to analyze these possibilities is by segmenting the network according to service areas.

Network Segmentation

To identify the problem more quickly, it is helpful for the administrator to keep the network topology in mind.

An effective way to address slowness issues from a topology perspective is to divide the network into service segments. This document proposes a division into four parts, which are common to most topologies:

- **Access Network:** the local network where the client’s PC is connected. This can be a wired or Wifi network and extends up to the first router.
- **Distribution Network:** the segment that connects the access network to the outbound network leading to the Internet. It may include one or more routers, depending on the internal topology.
- **Outbound Network:** the segment directly connected to the provider's router/modem.
- **External Network:** the Internet itself, starting from the provider’s network.

Figure 3 helps to better understand this segmentation.

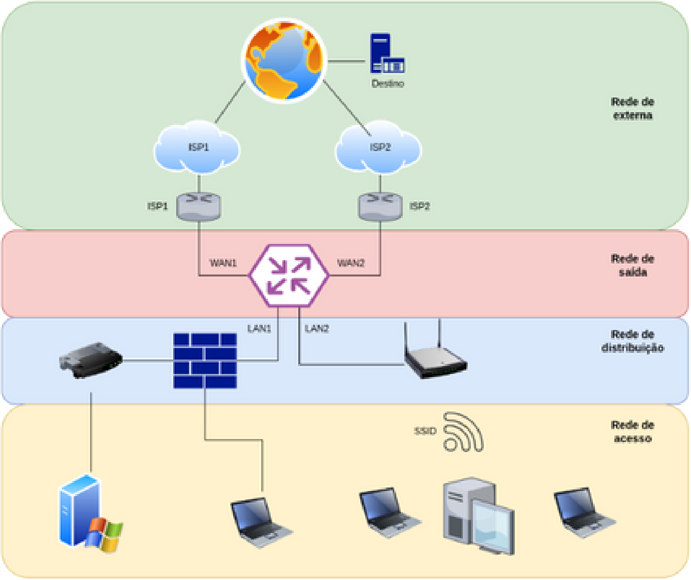

Figure 3 – Segmented Architecture to Facilitate Troubleshooting

Access Network

1.Ethernet Network

Ethernet has clearly dominated the access environment, and for good reason. The performance of Ethernet is largely due to the simplicity of its medium access protocol (CSMA-CD), which listens to the bus before transmitting a packet. Due to its performance and low cost, Ethernet has evolved rapidly in terms of speed. We've gone from the old 10 Mbps half-duplex networks to 10 Gbps full-duplex, with many variations in between. However, despite these advantages, Ethernet performance quickly degrades with cable length. For example, a 1 Gbps Ethernet network has a maximum length of 100 meters. Beyond that, performance drops from 90% to around 70%.

2. Wifi Network

Wifi began to scale significantly around 1999 with the launch of the 802.11b standard, reaching a maximum of 11 Mbps. Since then, several improved standards have emerged, currently reaching Wifi 6, capable of up to 9.6 Gbps.

Wifi solves various infrastructure issues by not requiring cables between the Access Point (AP) and the users. This access is provided via radio signals in the open 2.4 GHz and 5.8 GHz bands. These benefits and the good performance of Wifi networks—around 80%—have made their use explode in Brazil and worldwide.

For example, if you live in a large apartment complex and scan for Wifi networks with your phone, you'll likely find more than ten SSIDs being broadcast. Naturally, there is a cost to this level of usage, and that cost is interference.

Since Wifi uses open frequency bands, all access points share the same frequencies. So, when all your neighbors are using their Wifi networks, you'll notice a drop in performance and increased slowness.

Another issue with 5.8 GHz Wifi communication is that higher frequencies are more affected by physical obstacles. In high-frequency ranges, any obstruction has a larger impact and worsens signal quality. Therefore, maintaining good signal strength on 5.8 GHz requires installing a larger number of Access Points (APs).

In summary, despite its convenience, Wifi communication is not as efficient as a wired structure. Thus, a slowness complaint from a Wifi user may well be related to the access network.

Distribution Network

It is very common to find internal segregated networks used to serve different purposes. For example, companies that serve the general public often have a separate Wifi network for public use. This public network should be segregated for security reasons, and a good way to ensure this is by using a router to separate it from the company's corporate network.

Segregated networks that depend on internal routers or firewalls to reach the Internet often suffer when these devices have performance issues. Small and inexpensive routers and firewalls can be attractive and sufficient in most cases, but as the network or the number of users grows, things can become problematic.

A good way to detect such issues is with the **tracert tool** (on Windows). Running **tracert** from the complaining user's PC will show the network segments the packet passes through on its way to the Internet and the time taken across each segment. Comparing **tracert** logs from inside and outside the user's network can be very helpful in diagnosing and resolving the problem.

Outbound Network

The outbound network receives all traffic generated internally and routes it to the external network through the provider's equipment. If there are multiple providers, the outbound network is also responsible for implementing **failover** and/or **traffic balancing**.

1.Failover

Functionality that detects connectivity failure to a provider and activates redundancy.

2. Traffic balancing

Functionality that balances traffic entering the LAN among the available WANs. When implemented intelligently, traffic balancing helps with bandwidth aggregation.

The best way to diagnose slowness involving the outbound network is by running **tracert tests** from different sources to the same destination. Comparing recent **tracert logs** during failure events with older logs from normal operation is extremely useful for diagnosis and problem resolution.

External Network

The external network comprises the provider's modem/router and their infrastructure up to the Internet. Common issues like fiber cuts or failures in provider core or edge routers can degrade outbound network performance, which affects the access networks as a whole and impacts user experience —leading to slowness complaints.

Again, comparing old and new **tracert logs** greatly aids in identifying the problem and diagnosing the cause.

More modern and capable equipment on the outbound network can detect these provider-side performance issues and perform intelligent load balancing by routing more traffic through the better-performing provider, in cases where multiple providers are in use.

Conclusion

Regardless of the method proposed in this document, the first thing a network administrator should do when handling a slowness complaint is to conduct a preliminary review of available alarms and dashboards before interviewing the complaining user.

This simple step makes the conversation with the user much more productive, as it guides the testing based on the alarms observed. Keep in mind that users are not expected to have the technical knowledge to assist in testing.

Slowness issues can have many causes, which may even occur simultaneously, making diagnosis more complex. For this reason, developing clear methods and keeping historical data is vital for effective diagnosis and resolution.

TELCOWEB

Learn more about smart corporate networks!

Click the button below to get more information and personalized quotes!

I want more information!

telcoweb.com.br