



Connection Process

WHITEPAPER



Table of Contents

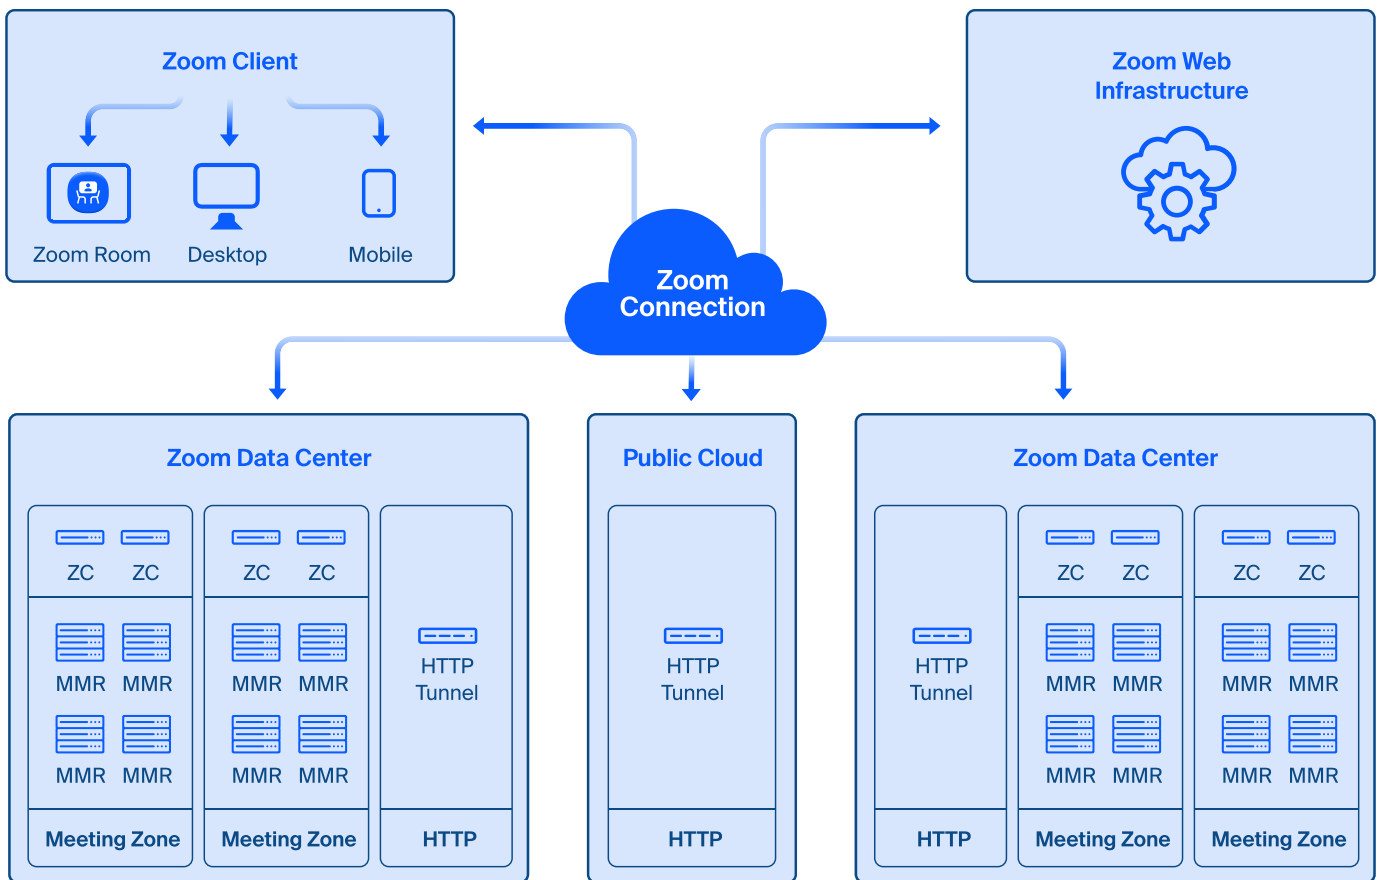
03	Overview
03	Core Concepts and Components
05	Connection Process Flow
05	Special Cases
06	Conclusion

Overview

Zoom is the leader in modern enterprise video communications, with an easy, reliable cloud platform for video and audio conferencing, collaboration, chat, and webinars across mobile devices, desktop computers, telephones, and room systems. One of the key differentiators that facilitates the ease and reliability of the cloud platform is Zoom's connection process. Zoom's connection process ensures that whenever someone attempts to access the platform there is an optimized path to Zoom's geographically distributed and highly available infrastructure. This white paper discusses that process and the technology behind it.

Core Concepts and Components

Prior to diving into the process, it's important to understand the key components involved within the connection flow and their role with Zoom's architecture.



Zoom Client

The Zoom Client is an individual's primary method for accessing the Zoom cloud. While available for multiple operating systems (macOS, Windows, Linux, Android, iOS, Chrome OS) and in a range of context-aware applications (mobile, desktop, Zoom Rooms, SDK), its interaction pattern with the Zoom cloud remains the same across all configurations.

Zoom Web Infrastructure

The Web Infrastructure is a highly available web application that not only helps host the zoom.us website accessed by many individuals every day, but also helps service application requests through its extensive API resources that are leveraged by external developers and the various components of the Zoom infrastructure.

Zoom Meeting Zone

A Zoom Meeting Zone is a logical association of servers that are typically physically co-located that can host a Zoom session. A Zoom Meeting Zone and its associated servers may be located within one of Zoom's global data centers or can be located within an organization's network if running Zoom's on-premise solution. The primary components of a Meeting Zone are Multimedia Routers and Zone Controllers.

Zoom Zone Controller

A Zoom Zone Controller is responsible for the management and orchestration of all activity that occurs within a given Zoom Meeting Zone. Deployed in a highly available configuration, these systems track the load on all servers with the Zone and help broker requests for new connections into the zone.

Zoom Multimedia Router (MMR)

A Zoom Multimedia Router is responsible for hosting Zoom meetings and webinars. As the name implies, these servers ensure that the rich offering of voice, video, and content are properly distributed between all participants in a given session.

Zoom HTTP Tunnel (HT)

The Zoom HTTP Tunnel service is an integral part of Zoom's network resiliency strategy. Housed in various public clouds and Zoom data centers, these servers offer a connection point to clients who are unable to connect to the Zoom platform through other network channels. Once a tunnel is established between the Zoom Client and Zoom HTTP Tunnel, the client is able to access the Zoom Meeting Zone across the various data centers.

Connection Process Flow

The process of connecting to Zoom session is divided into four phases as outlined below.

Meeting Lookup

Upon receiving a request to join a given session, the first action taken by the Zoom Client is to contact the Zoom Web Infrastructure to obtain the applicable metadata required to access the meeting or webinar. Accomplished over a HTTPS connection using port 443, the Zoom Client uses this opportunity to better understand its current network environment including details such as proxy server usage. On the other side of the connection, the Zoom Web Infrastructure prepares a package of data optimized for that client. Through the use of Geo-IP and other Zoom service delivery technology, a list of optimum available Zoom Meeting Zones and associated Zoom Zone Controllers are returned to the client along with meeting details so it can proceed to the next phase in the connection process.

Meeting Zone Selection

With a list of Zoom Meeting Zones that could service the Zoom Client for the session, the connection process then enters the next phase of the workflow. To ensure the best connection is used, the Zoom Client attempts to connect to each of the Zoom Zone Controllers within the Zoom Meeting Zones provided in the previous phase and then conducts a network performance test. By comparing these results, the client is able to confirm there is a connectivity path in place to each Zoom Meeting Zone and select whichever one demonstrates the best performance. Zoom's innovative protocol leverages HTTPS. This connection is attempted over SSL (port 443).

MMR Selection

With the ideal Zoom Meeting Zone selection from the previous phase, the client then requests details of the best Zoom Multimedia Router (MMR) from the Zoom Zone Controller. Once identified, the Zoom Client reaches out to the MMR directly to establish a control channel for the session. This connection leverages a protocol developed by Zoom which communicates via SSL on port 443.

Media Routing

With a successful connection to the optimum Zoom Multimedia Router for the session, the Zoom Client prioritizes creating a connection for each type of media that will be exchanged such as video, audio, and content. Each of these media connections attempt to use Zoom's own protocol and connect via UDP on port 8801. If that connection can not be established, Zoom will also try connecting using TCP on port 8801, followed by SSL (port 443). By leveraging different connections for each type of media, further network optimization technology can be applied such as DSCP marking to ensure the most important media is expedited through the network.

Special Cases

While the process outlined above covers most use cases, there are a few special exceptions that have been implemented to help ensure a reliable session even in complex networks.

Proxy Servers

During the Meeting Lookup phase of the connection process flow, the Zoom Client can determine if a proxy server is used as part of the network connection path. If one is detected, during the Meeting Zone Selection and MMR Selection phase of the connection process the Zoom client will immediately leverage the proxy server and attempt to make the associated connections to the Zoom Zone Controller and Zoom Multimedia Router using SSL.

HTTP Tunnel

If there is no response from any of the Zone Controllers after 5.5 seconds, the Zoom client will attempt to connect using the HTTP Tunnel. To ensure multiple paths for a successful connection, these servers are housed in both public clouds and Zoom data centers. This connection is attempted over SSL (port 443). The Zoom client will ping multiple HTTP tunnels and the first to respond is used.

Web Client

If the Zoom Client is unable to connect through any of the methods listed above, it will direct the user to connect to the meeting via the Zoom Web Client in their browser, without downloading any plugins or software. The Zoom Web Client and Web SDKs attempt to connect on 443 secure WebSocket for signaling. For Video and Audio media we attempt to connect to UDP ports 8801, 8802, 8803, and 8804, followed by TCP Port 8801 and 8802, and finally followed by TLS port 443.

Conclusion

A growing number of businesses, small and large, rely on Zoom services every day. Zoom offers multiple connection paths utilizing various protocols across a geographically distributed infrastructure to ensure a successful connection for all users.