

Mail från: Jane Ross, Zoom, till markus.schneider at sunet se 12 mars 2020

I appreciate your concerns about Zoom's capacity and elasticity in light of the current situation. As you know, we are fully committed to doing everything we can to provide resources and support those impacted by the coronavirus outbreak. We recognize that with this commitment, it is critical to ensure our architecture can withstand increased use.

Foundationally, we have provided a video first architecture from day one that is designed to scale. There are a few things you should know about how we operate our infrastructure and our business that we expect will give you confidence in our ability to serve our global customers amidst increasing demand. Specifically:

1. We operate our own global (co-location) data centers in 17 locations around the globe, providing significant control and flexibility when it comes to routing both our audio and our video traffic.
2. We have always ensured that we have enough capacity to handle double our average daily peak. This has always been our policy, even before this global health crisis.
3. The resilience of the Zoom platform has also been key to us from day one. Accordingly, we have architected the platform such that, in the event of capacity constraints at the data center nearest a user, additional traffic will be routed to one of our other data centers.
4. In the case of an unprecedented, massive influx of demand, we have the ability to access and deploy tens of thousands of AWS servers within hours, and in full compliance with our Privacy Policy (zoom.us/privacy), to seamlessly scale without any impact to our users.
5. One area where we may have some potential constraint is related to PSTN (phone numbers) used for people to dial into meetings. [In rare instances,] the local carriers that operate those dial in numbers hit their limits in some [regions] during periods of extremely high usage. In those situations, VOIP audio can be leveraged through Zoom to communicate seamlessly without having to rely on the carriers. We would suggest considering using mobile phones in this instance. While this is an inconvenience, note that it has no impact on our ability to continue hosting meetings.

In case of interest, we have also added below an excerpt below from our S-1 filing last year that walks through some other unique aspects of Zoom. As you'll see, these characteristics have enabled us to build meaningful scale to manage situations like this.

Please let me know if you have any other questions.



<https://www.sec.gov/Archives/edgar/data/1585521/000119312519083351/d642624ds1.htm>

Our Technology and Infrastructure

Our unique technology and infrastructure enable best-in-class reliability, scalability and performance. We designed our communications platform to be video-first and cloud-native. Most legacy approaches utilize single multipoint control units (MCUs) to bridge video and voice participants into an integrated stream that is broadcast back to the participants. These hardware devices are shipped with defined processing and memory capacity that are difficult to scale. In addition, an MCU architecture is similar to other mainframe-like approaches where stream processing and mixing run on the same machine, which is resource-intensive and limits scalability.

Our technology was specifically designed from the start to address the most difficult component of communications: video. Video requires intense computing resources for encoding, decoding, multiplexing and synchronization, as well as higher bandwidth and network performance, to a much higher degree than other forms of communication like voice, chat and content sharing. Our architecture separates video content processing from the transporting and mixing of streams. We allocate video content processing to intelligent agents that reside on client devices and dynamically encode and decode based upon the performance of client technology, network performance and bandwidth. We leverage a next-generation multimedia router (MMR) that operates on commodity hardware and a globally distributed cloud infrastructure to determine the optimal data centers to host a meeting and an optimal set of paths to connect the participants.

Our technology and infrastructure have the following key strengths:

- Architectural separation of content processing from stream routing.** Client-side processing for encoding and decoding, our MMR for routing and dynamic optimization of stream traffic enables us to provide reliable high-quality communications. This architecture means each MMR can support approximately 1,100 more meeting participants than a standard MCU, which generally supports up to 80 participants.

- ***Distributed cloud-native infrastructure from the start.*** From day one, we have invested in a distributed architecture and load balancing to ensure that end users always get the best possible performance and quality. Our geographically distributed architecture enables users to connect to the data center closest to their location. Our distributed architecture also enables us to provide a more scalable service. Our microservices architecture allows us to access resources automatically that enable us to seamlessly grow our capacity. As a result, from a technology perspective, we provide our users with a reliable video experience from two to thousands of video participants.
- ***Proactive quality of service application layer at client.*** Our optimization technology is based on proprietary algorithms that detect packet loss, latency, jitter, CPU utilization and bitrate/bandwidth to optimize the video, audio and content sharing experience. Depending on the type of device, our algorithm prioritizes the factors most important for the particular device, resulting in the best possible user experience.
- ***Security and disaster recovery.*** We offer robust security capabilities, including end-to-end encryption, secure login, administrative controls and role-based access controls. Our platform is SOC2 Type II and TRUSTArc compliant and enables customer HIPAA compliance. We have resilient disaster recovery capabilities, including active-active disaster recovery for real-time processes to help ensure that backup active servers are available.

Our cloud-native infrastructure offers the following key strengths:

- ***Private cloud for real-time activities.*** We host real-time video and voice using 17 co-located data centers on our own servers around the world. We have built a network of private links so that each data center connects to multiple others. This structure creates redundancies and ensures reliability and performance.
- ***Optimized public cloud usage.*** For maximum efficiency, we deploy public cloud providers (Amazon Web Services and Microsoft Azure) to host our web applications and messaging service that do not require real time usage.