

GOALS

Each of our customers have different usage profiles of their Video-Rooms, some have one to one meetings, some have small group meetings and schools have up to 100 users in one room. From our current experience, it is difficult to define the needed server infrastructure to allow the flexibility of providing a high available service in all those scenarios, in a cost effective way.

In our experiment we have tested and measured the capacity of our SFU in the different scenarios, to ensure appropriate server configuration for various customers

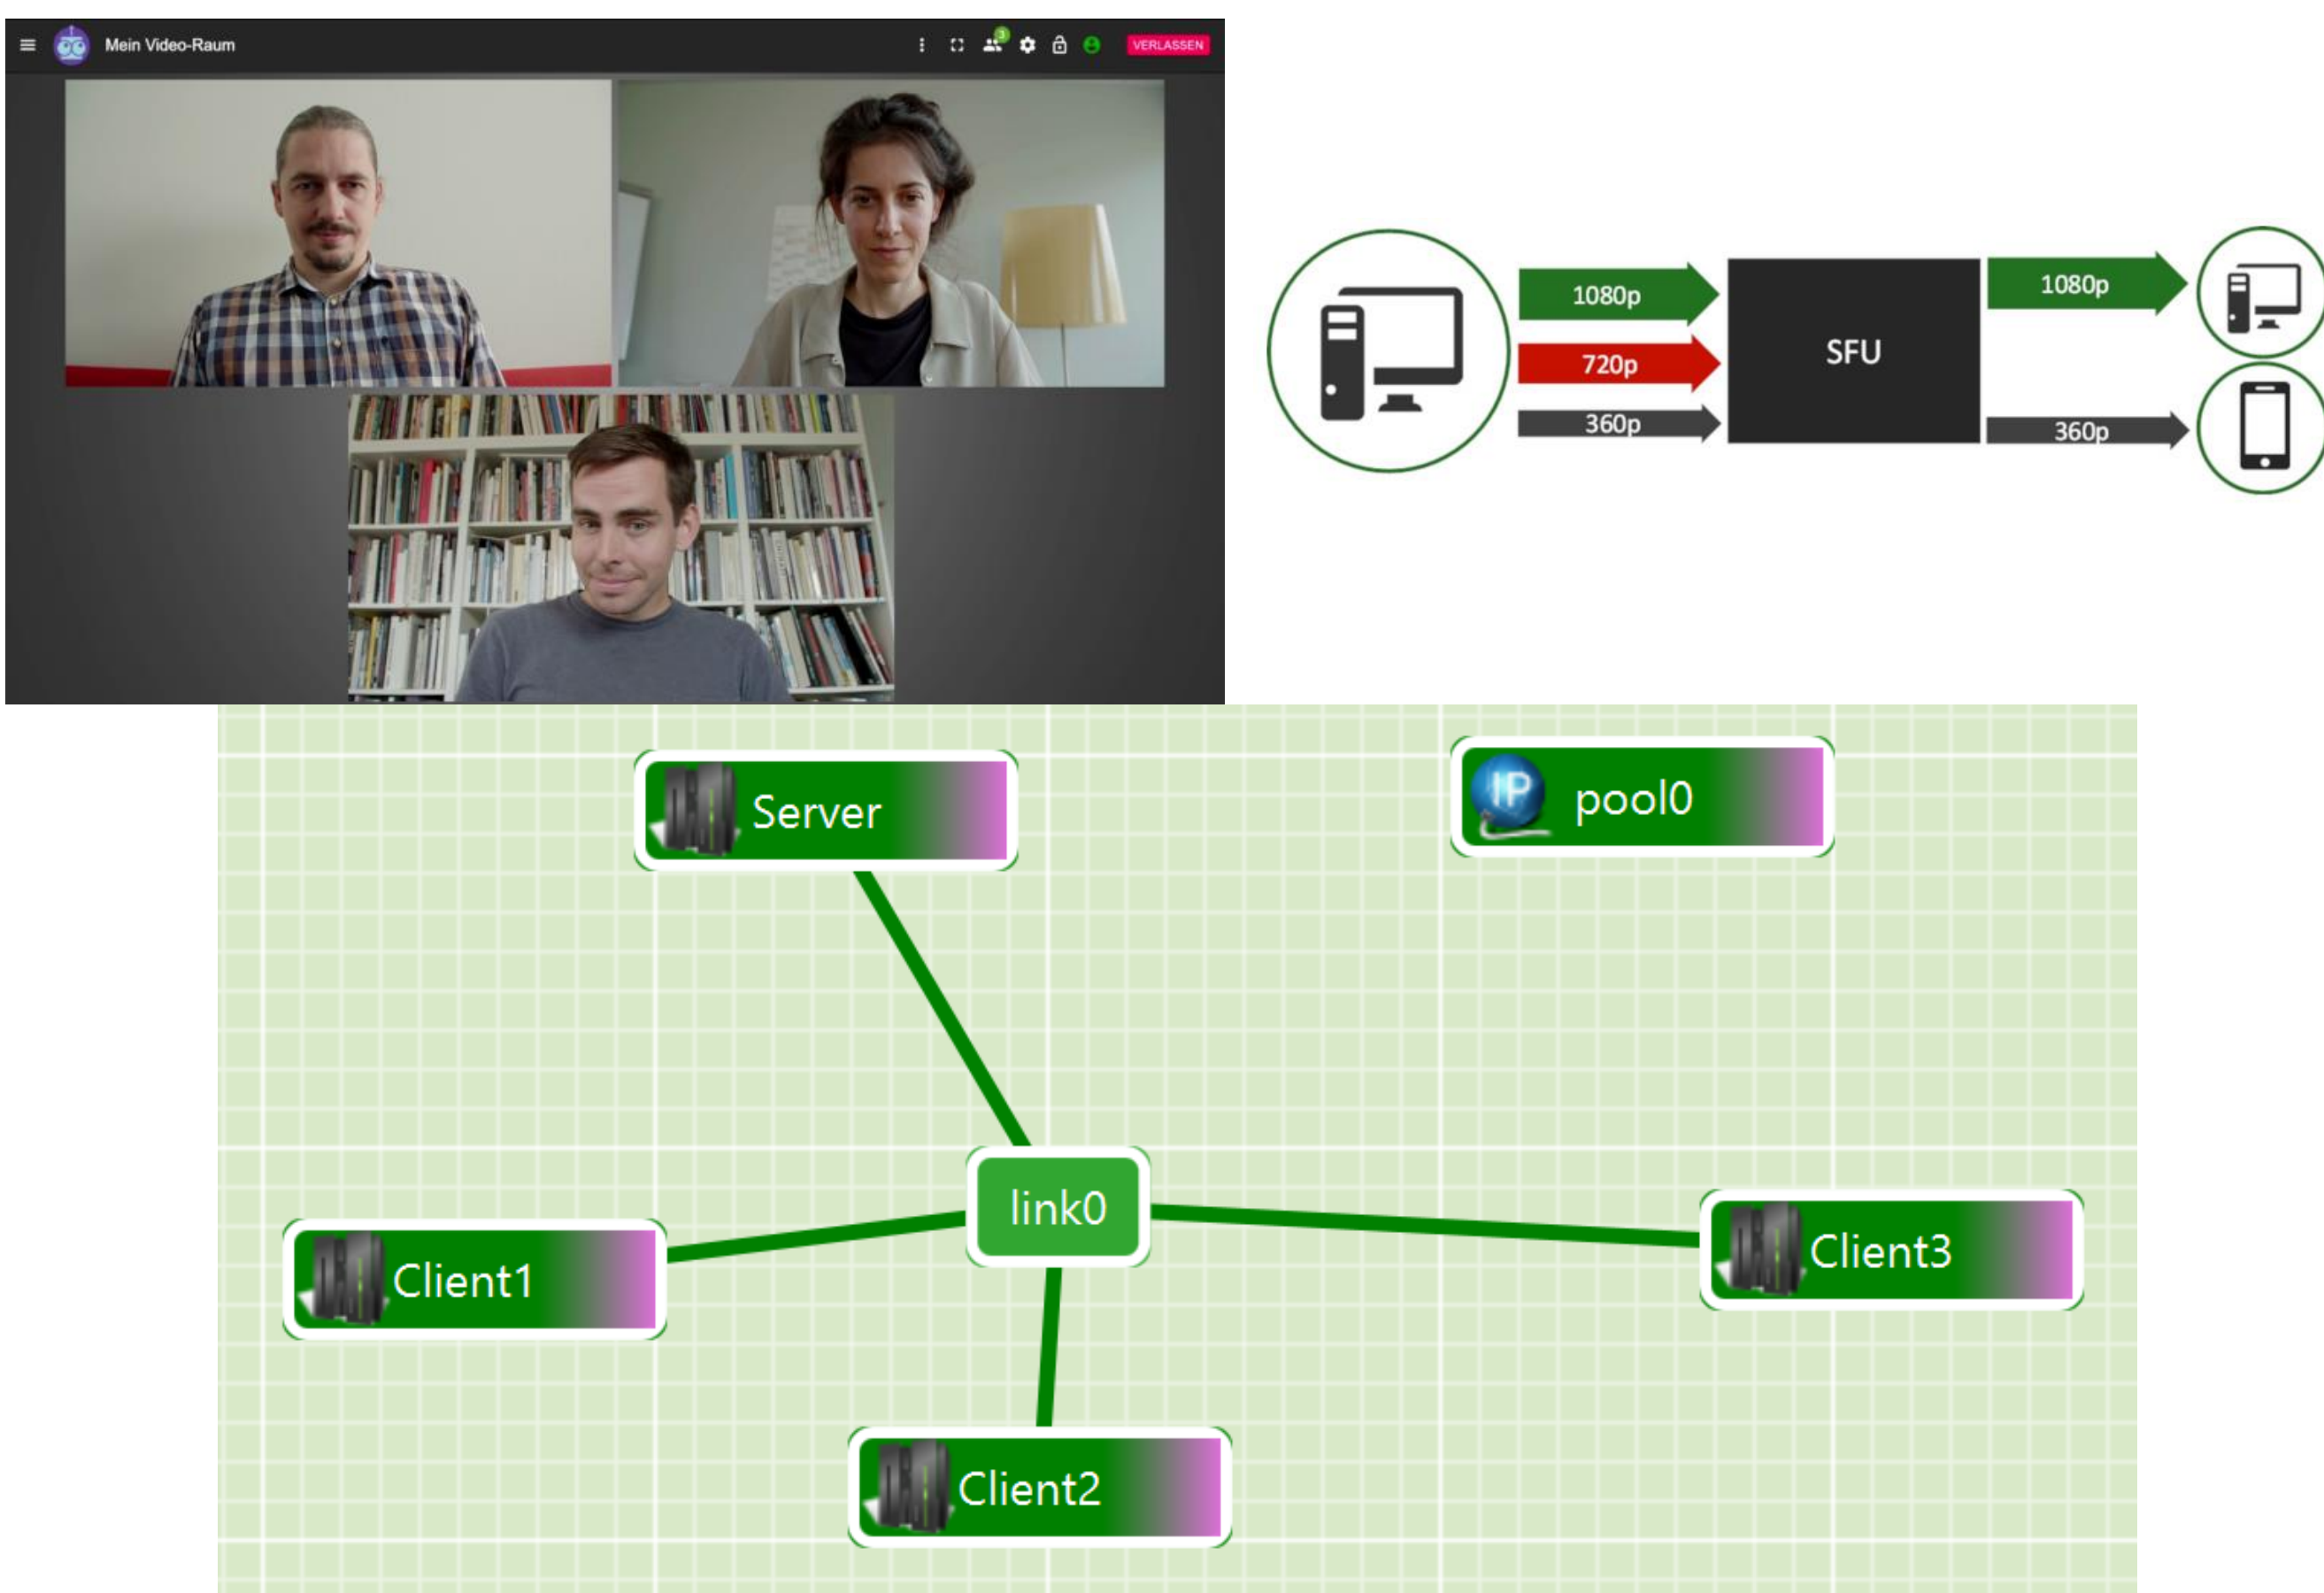
CHALLENGES

The load per CPU and the total load of the machine depends on such parameters as:

- number of users in a room
- number of enabled video producers (webcam, screen sharing, additional video inputs)
- number of enabled audio producers (microphone, screen sharing audio)
- number of active video consumers (users having a given video displayed)
- number of active audio consumers (users listening to a given audio stream)


We have defined a set of most common usage scenarios, varying the parameters listed above, and run multiple test combining rooms of same and various scenarios on a server.

DEMO SETUP



RESULTS

Max. users in 1 room with router size 40



MORE RESULTS

We have run experimentation with grouped scenarios of the same type. This has allowed us to measure the maximum rooms that a server with a given number of CPU cores can handle and the maximum number of users per router.

We have run the experimentation for the following grouped scenarios:

- 2 users in 1 room on 1 worker – how many rooms per worker/CPU core supported?
- 4 users in 1 room on 1 worker – how many rooms per worker/CPU core supported?
- 20 users in 1 room on 1 worker – how many rooms per worker/CPU core supported?
- 100 users in 1 room on multiple workers with router size 40 – how many workers/CPU cores needed?
- 1 room on multiple workers/CPU cores with router size 40 – how many users can access?

CONCLUSIONS

During our experimentation we have found out, that on a sever equipped with 2.40GHz CPUs, we can:

- Run 30 2-user rooms per CPU core
- Run 22 4-user rooms per CPU core
- Run 3 20-user rooms per CPU core
- Need 3 CPU cores for a 100-user room
- Can have a maximum of 240 user in 1 room

POST MORTEM

- Form the experiment we have gained new knowledge about how to better optimize our infrastructure usage.
- We have calculated safe numbers of various rooms that we can provide using a single CPU core
- The tricky part is that when adding an additional router on a new worker, you have an overhead of routing users between the workers.
- To enable better optimization of infrastructure usage, we have to group our users depending on the type of the meetings they conduct.