



FACTSHEET



Solution Scalability

Solutions that scale

The Lumi AGM Mobile platform is built on a micro services architecture that allows for resources to be added or removed on an ad-hoc basis. This provides the ability to scale the system up or down, dependent upon current needs and customer requirements. Additionally, it is built in a manner which allows completely standalone server configurations to be setup for specific meetings or customers to ensure scalability and integrity. This is primarily done for regionality (e.g. allowing meetings in certain geographical regions to connect to local servers). The current configuration includes three hosting zones (Europe, Australia and the US).

The architecture is further fragmented into separate 'responsibilities' for servers, allowing specific heavy tasks to initiate additional nodes to cope with the temporary strain of large-scale meetings (i.e. based on number of participants connecting/joining) or meetings that may require heavier interaction with connected participants.

In Production Load Balancing

The production system uses Amazon Web Services (AWS) load balancers. These are capable of spreading the load to the available nodes within the regional cluster that is in use. Once the connection with a node is established, a participant's login credentials are certified and they are granted access to the meeting if their credentials are valid (i.e. a valid Meeting ID followed by correct username and password). The participant database (i.e. customer data) is separated from the internet and only Lumi's AGM servers can query them directly. This ensures that any potential malicious attacks cannot access the data directly.



Load Tests

Before every major release, the development team performs a barrage of load tests to simulate different types of meetings, meeting behaviours and edge cases that might be seen during production. A load test client is employed that uses the HTTP API to connect and this allows for automation of all behaviours that might occur in production. Additionally, it also allows for the easy scripting of new scenarios that represent different usage types. Further to the automated tests, a mobile application is used to connect directly to the meetings to get a subjective view of how the system is working and when it starts demonstrating slow or unresponsive behaviour.



Solutions that scale

The scale testing performed shows that meeting size is more important than total connections to the servers. This means that, for example, one meeting of 10,000 participants is much more consuming of system resources than 10 concurrent meetings of 1,000 participants.

The system has no issue in coping with a 10,000 participant meeting in its base setup and can scale to run larger meetings by adding additional nodes for processing performance. The development team recommends that where possible, these nodes are added prior to the meeting instead of using automatic scaling in the middle of the event. However, please note that this is only a precaution as automatic scaling has been seen to repeatedly work successfully during testing.

The different functionalities of the solution take up differing amounts of resources. To date, the most intense task seen is the registration of participants, therefore it is recommended that larger meetings leave enough time to allow all participants to register before the first vote is pushed out. Recommended time is 5 minutes per 5,000 participants.

As an example, the current base configuration can easily facilitate 15 meetings of 1,000 participants, 30 meetings of 500 participants or 50 meetings of 200 participants, at the same time. It is also possible to run one meeting of 10,000 participants and multiple sub-1,000 participant meetings at one time.

All meetings which utilise the Lumi AGM Mobile platform are scheduled in advance, giving Lumi complete visibility of all meetings taking place, globally, and on each day. The system architects analyse the size of each meeting on a regular basis and should there be a period of time when the system might be under a greater load (in terms of connections) than normal, contingency measures are enforced and configurations will be made to increase the processing power within a specific server.

Adding additional processing power primarily boosts additional participant numbers whereas more powerful nodes will help serve larger meetings. Lumi therefore recommends larger machine instances if a meeting will be of significant scale (i.e. over 10,000 participants).