

# Case Study



## Systems Integrator

---

### Background

A systems integrator was bidding for a large multi-million pound health authority project, which involved a requirement for a sub-second average response time. The systems integrator recognised that the target was an unrealistic response time for a multi-platform architecture, but they didn't know what response time they could commit to. As the health authority had set severe penalties if these performance targets were not met, the systems integrator needed to determine the achievable response times quickly before it committing to them.

The systems integrator recognised the need for performance analysis of the proposed architecture. Capacitas was contracted to analyse the existing situation and produce response time estimates.

### Approach

Capacitas undertook an analysis of the proposed architecture and applications, and conducted a series of discussions with the architects in order to identify the most appropriate approach to determine expected response times.

Capacitas proposed the building of a queuing model of the system using a common desktop analysis application. The analysis was used to model the end-to-end delay for each key transaction type across the architecture. This included modelling the application, system and network layer.

This model used analytical queuing, as this method can be developed within a relatively short period of time and also has a good degree of accuracy. Capacitas took this approach, as the systems integrator needed answers within the period of a month due to the bidding timetable.

## Problems

The following were a few of the problems identified in creating a meaningful queuing model:

- The architecture was not well defined
- Transaction flows through the architecture were undefined
- The transactions were in a Health Level 7 (HL7) message format using XML, and the parsing times for these transactions were unknown
- The sizing of the architecture was in a state of flux
- The authority's requirement for service level agreements was unclear
- The mapping of different response time requirements to different transaction types was not defined

## Impacts

- The proposed architecture was over-sized. The peak hour utilisation figures were less than 10%
- The systems integrator's architecture team was under pressure to meet the authority's response time requirement. The architecture team didn't have any evidence to show that the authority's response time requirement was not achievable

## Our Solution

The results from the Capacitas queuing model found:

- That the transactional response time requirement was not achievable using the existing architecture
- That depending on the transaction type, the actual transaction response times would range from 2 times the target response time for an optimally tuned system to 4 times the target for a poorly tuned system
- That the architecture was oversized and that this excess capacity would have little or no improvement in response times
- This information was used to significantly reduce the bid cost

## Benefits

The benefits of the analysis were:

- The systems integrator now had estimated values for

transaction response times. Capacitas had provided the systems integrator with expected response time figures, which supported the case for the systems integrator not to attempt to meet the authority's response time requirements

- The analysis dispelled assumptions that the systems integrators architecture was unable to meet the response times because of network delays. Capacitas identified that the major delay was occurring at the application layer rather than the network layer. The network delay contributed less than 5% to the overall delay. This helped support the systems integrator's bid, as the authority believed that the architecture would have excessive delay due to the network components
- Analysis showed that adding additional CPUs to the system would not necessarily improve response times
- A model was produced that was subsequently used for costing of the project bid. As the authority wanted the costs for different response time requirements Capacitas used the model to cost these different response time options

## Conclusion

The systems integrator went on to win the bid, even though rival bidders had agreed to meet the authority's response time requirements.

The Capacitas queuing model was used as evidence that the authority's response time requirements were not achievable on a multi-platform architecture and the authority accepted this evidence.

© Capacitas Limited 2002.

Capacitas  
Prince Consort House  
109-111 Farringdon Road  
London EC1R 3BW  
United Kingdom