

Report on the performance of the 2nd floor J block labs

- **full report** `<http://floppsie.comp.glam.ac.uk/Southwales/gaius/lab/performance/report.pdf>`

- **Problem**
 - reported 40 minute login times on Window clients in 2nd floor labs

- **Possible solution**
 - over the next year there is a possibility of a (£86K) network upgrade from 100Mbps to 1Gbps in the labs and from 1Gbps to 10Gbps to the servers
 - is this likely to deliver the benefits expected?
 - measure the network performance in the 2nd floor labs


The experiments

- first experiment used an artificial workload (infinite source/infinite sink)
 - each client wrote 64MB to the server using a range of block sizes

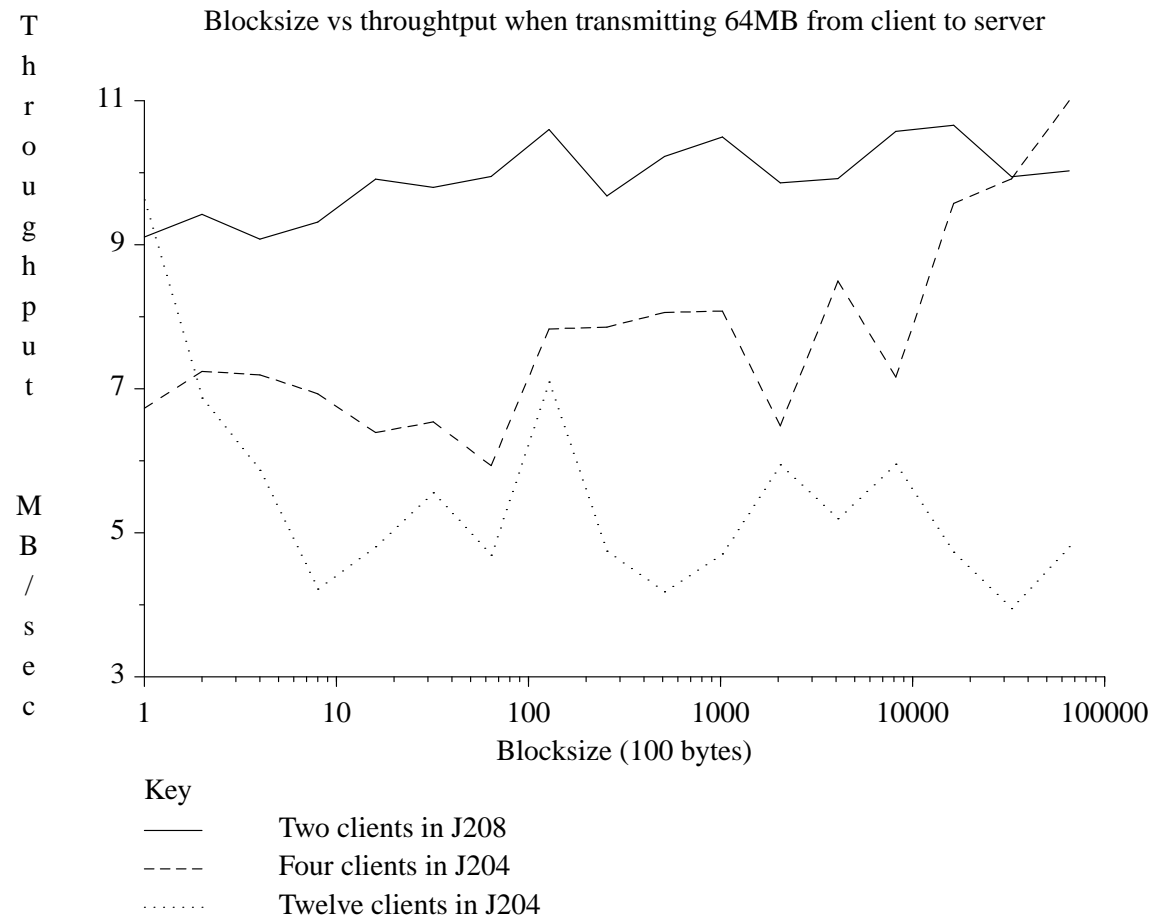
- second experiment, 2nd year coursework download (6.5MB)
 - each client downloaded a `.tar.gz` file and extracted the contents
 - the unpacked contents is written back to the server

- third experiment, 3rd year coursework download (137MB)
 - `ioquake.tar.gz`, extracted and written back to the server

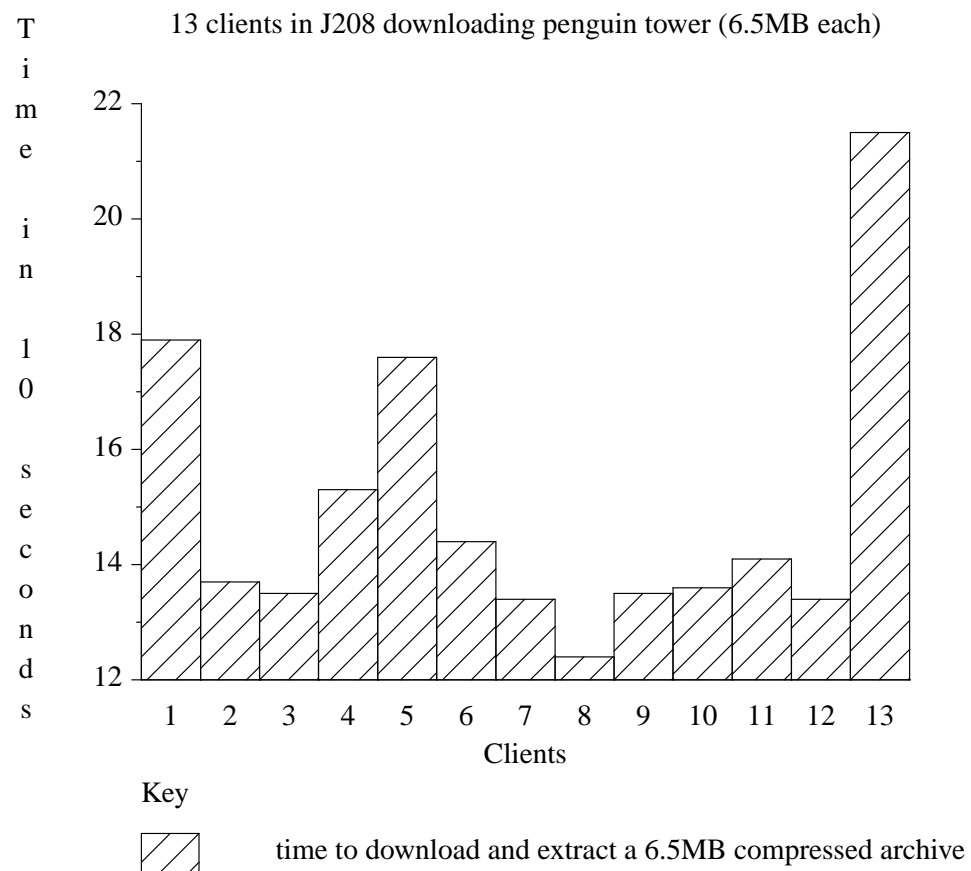
Network diagram



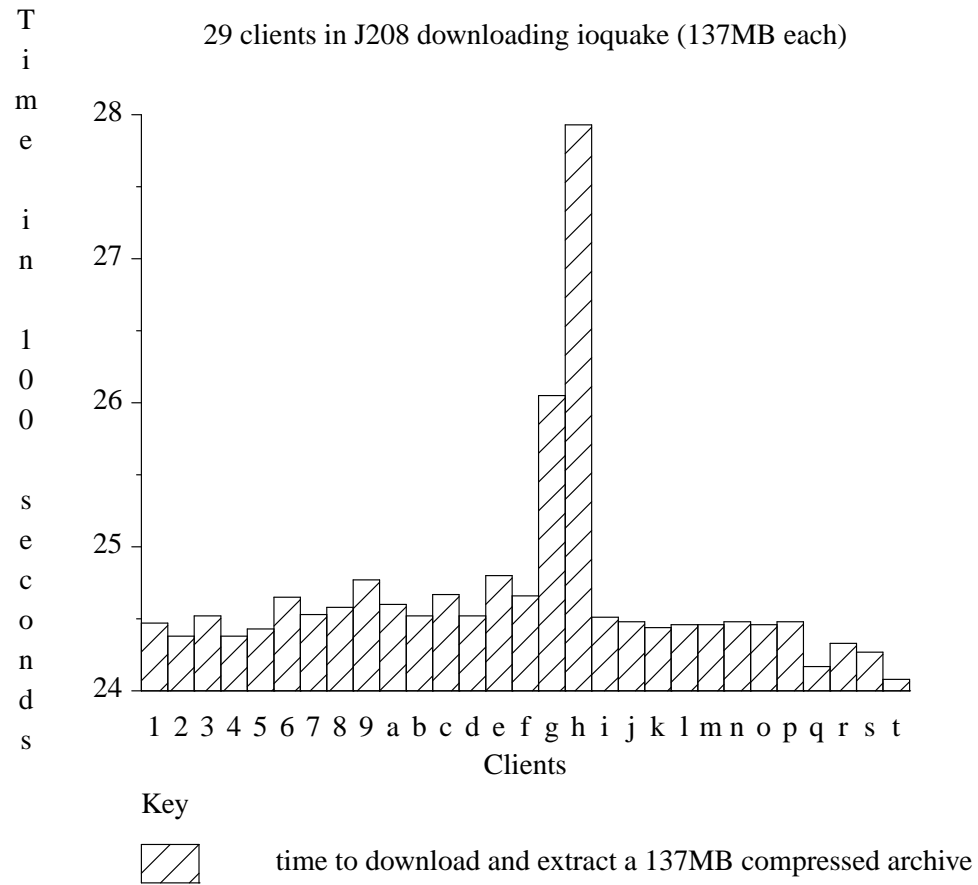
Results: experiment one



Results: experiment two



Results: experiment three



Summary

- the first experiment says that each client can receive 4MB/sec
 - up to a maximum of 12 machines
- the third experiment tells us that the network degrades evenly
 - 29 machines in J208 had pretty near the same elapsed time

Conclusion

- the network performs at the speed of the lowest link in the route
- unless the Windows patches exceed: $\sim 10\text{GB}$ ($4\text{MBs} \times 40 \times 60$)
 - it is unlikely that the network infrastructure upgrade will reduce login times
- future tests
 - substitute a router connecting the lab machines in J208 to the backbone with a GNU/Linux machine running a packet analysis tool `netwatch`

Conclusion

- Andrew Tanenbaum's[1], sometimes non obvious, rules for increased network performance which state in order of importance: “CPU speed is more important than network speed”, “reduce packet count” and “minimise copying”
- obvious hypothesis is that having a local high performance domain server (or reducing the data sent by the domain server) is likely to yield significant gains.