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

Gaius Mulley

ABSTRACT

This document describes an interim investigation into the network performance of the second floor J block laboratories. The reason for the report is that users are complaining about poor network performance in the labs. In particular there have been occasions when users have taken 40 minutes to login to the Windows client. The tests here were all conducted under the GNU/Linux utilising the wired network between the 2nd floor labs in J block and the departmental server room J105. The tests will determine whether there are any network ports which are dropping back to a lower throughput and also how the network infrastructure behaves under load.

Over the next year there is a possibility of a (£86K) network upgrade from 100Mbs to 1Gbps in the labs and from 1Gbps to 10Gbps to the servers. Clearly we need to know whether this is likely deliver the benefits expected. The report presents the performance findings, makes the case for further targeted experiments and provides some metrics for the likely improvement (or lack) of the Windows logon problem based on network upgrade.

This report attempts to measure the performance of the 2nd floor labs in J block. For all tests the laboratory was assumed to be otherwise idle (the tests were run in the evening or late Friday afternoon). There were three network workloads employed. The first, artificial, test scenario was an infinite source infinite sink application run between client and server. In each test run the client was required to write a 64 MB file to the server over a wide range of block sizes.

The second, real world, test scenario involved the clients downloading and extracting a coursework 1 for a second year games module. This coursework had been set the same week. The file was a 6.5MB MB compressed tar archive (penguin tower) which would be downloaded from the server to the client, unpacked on the client and the resultant unpacked data sent back to the server.

Lastly the third, real world, scenario involved the clients downloading coursework 1 for a final year module which was also issued during November. This archive was 137 MB (ioquake3, again compressed tar downloaded from the server, extracted and resultant files uploaded to the server).

Network infrastructure and the artificial test case 0

This section provides an overview of the network infrastructure tested. It is important to recognise the limits of the tests. The limitations are reasonable in the initial phase of network performance investigation, to restrict the possible interactions between competing application or operating system demands. Ideally it would be good to vary one element of the scenario and remeasure the performance make some interim conclusions and adjust the test scenario and repeat.

These tests were conducted on GNU/Linux clients (Mint 13) connecting to the GNU/Linux server mcgreg (running Debian Squeeze). While this says nothing about the performance of a Windows based system, it does give a set of accurate data points on how the network infrastructure performs under load between the labs and the departmental server room.

In broad terms the clients in the 2nd floor labs (J203, J202, J204, J208 and J210) are all connected via 100 Base T 802.3 wired network, which is constructed in a small hierarchy of switches. The switches have a series of n $\times$ 100 Mbps links and 1 $\times$ Gbps link. The server room has a series of 1Gbps links which feed into the top of the switches.

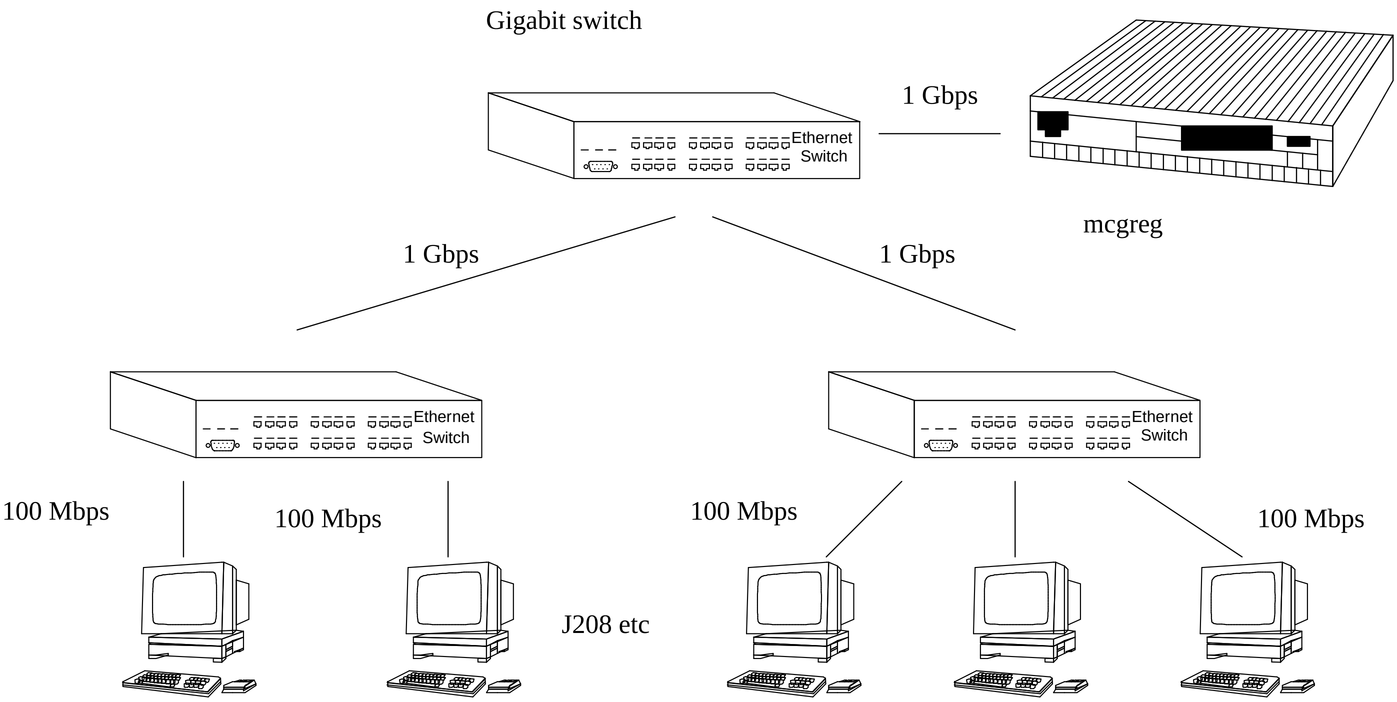

Figure 1: approximate network diagram

There were three independent tests run. The first test was to measure two clients simultaneously transmitting a file of 64 MB to the server. The second test, involved four clients transmitting 64 MB to the server and the final third test required twelve clients transmitting 64 MB to the server. The graphs give the performance of one of the clients in the experiment. The clients attempt to send a 64 MB file to the server using a range of block sizes. The operating system will attempt to utilise the most optimum block size (1500 bytes). The theoretical maximum throughput on an ideal network with computers running at infinite speed would be 12 Mbits/s.

The graph in figure 2 shows the performance characteristics of the three tests. It is worth noting that the network infrastructure (at the Physical layer) is in theory capable of achieving 100 Mb/s from the client and this is multiplexed into a 1Gb/s link into the GNU/Linux server (mcgreg). Reassuringly we see a linear performance increase as we add clients from 1 to 10.

It is interesting to see the performance drop off when we have 12 machines connected. Further work is required to explore the effect of 16 and 24 machines together with some real world loads.

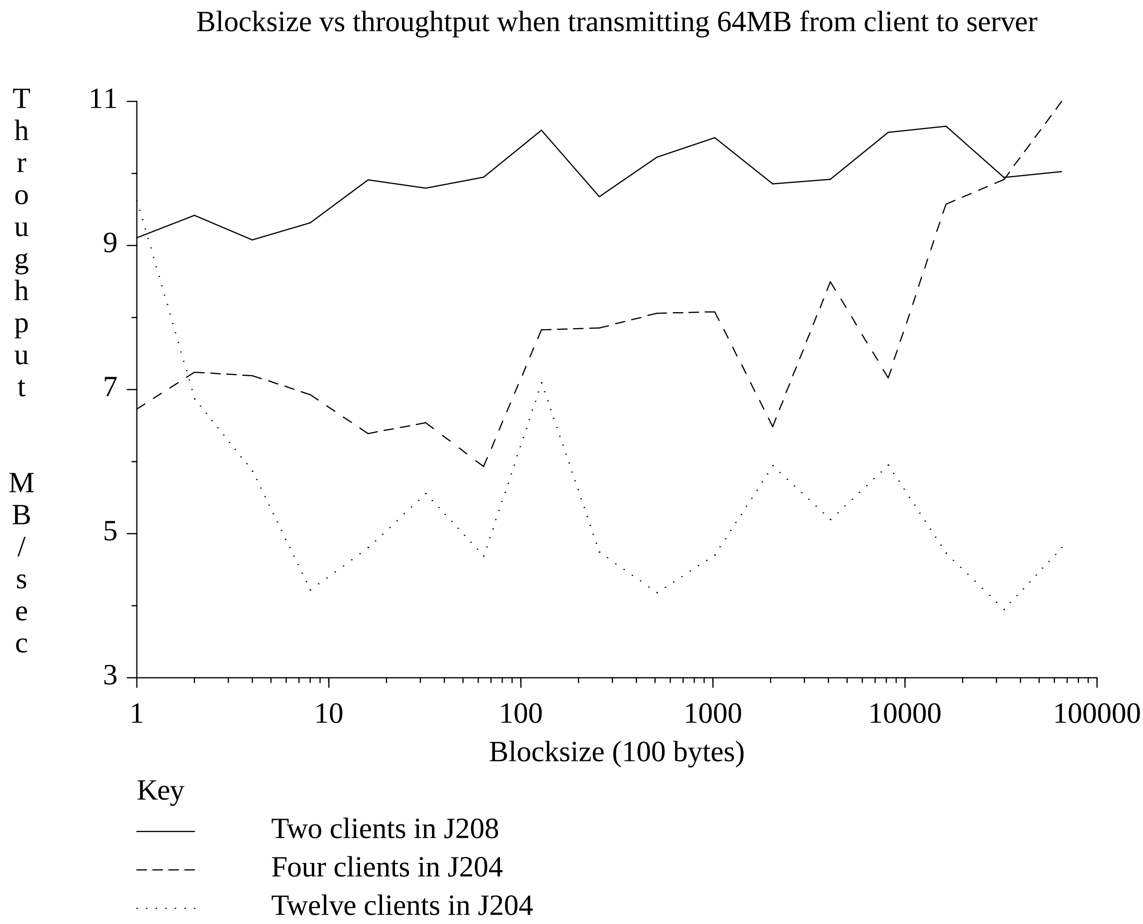

Figure 2: performance of infinite source infinite sink for 2, 4 and 12 clients

Test case 1: download of penguin-tower coursework

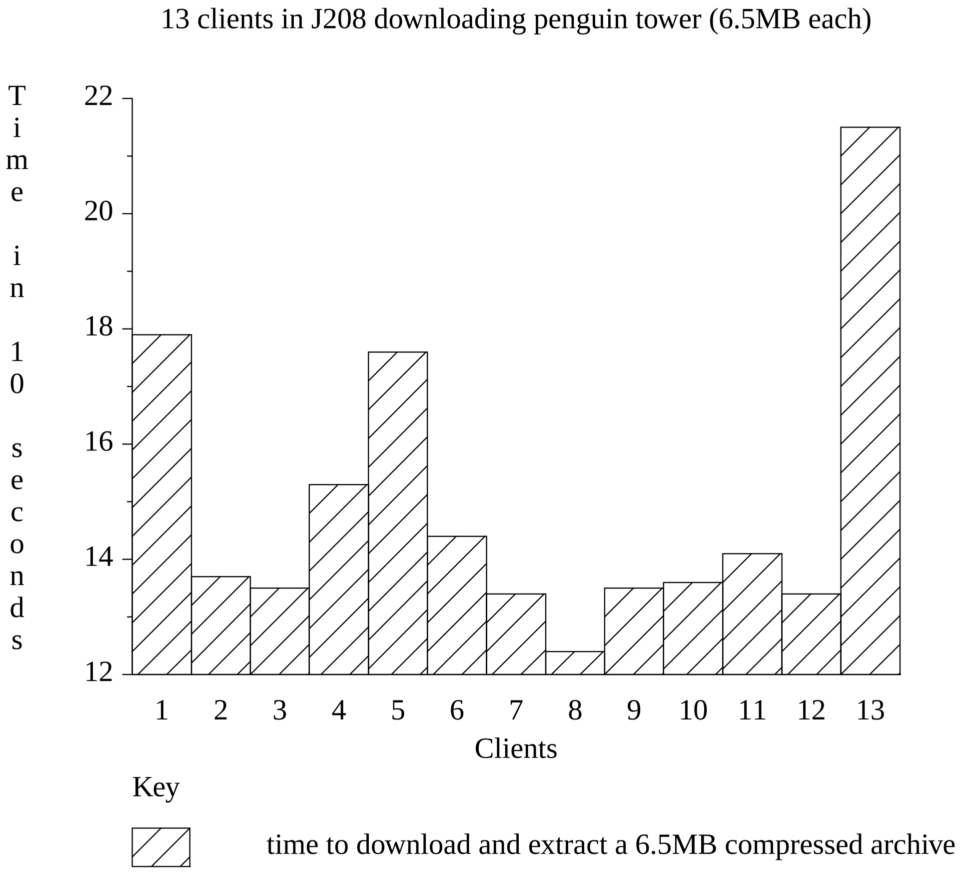

Figure 3: performance of 13 clients downloading and extracting 6.5MB archive

The next performance test conducted was a download of coursework 1 for second year games students. This archive was a 6.5MB tar compressed tar which was also uncompressed client side (the uncompressed data was automatically uploaded to the server). We see that 13 clients behave a little erratically, almost a 9 second difference between client 8 and client 13.

It maybe that the number of clients is too small or the size of archive is too small to gain any meaningful knowledge. The operating system will do everything to mask any inefficiency (by caching writes and using the most efficient block size transmission). It would be interesting to see how the network performs with a larger archive and many more clients.

Test case 2: download of ioquake coursework

In this test an archive of 137 MB is downloaded and extracted by 29 clients in J208. This coursework was released in November for the final year games students (an archive of the ioquake3 game engine with enhanced textures).

The bar graph of figure 4 shows the time ranged from between 40 and 46 minutes. It was encouraging to see that all clients finished within 16% of each other indicating confidence that the local departmental network scales well under load and distributes performance fairly across the clients.

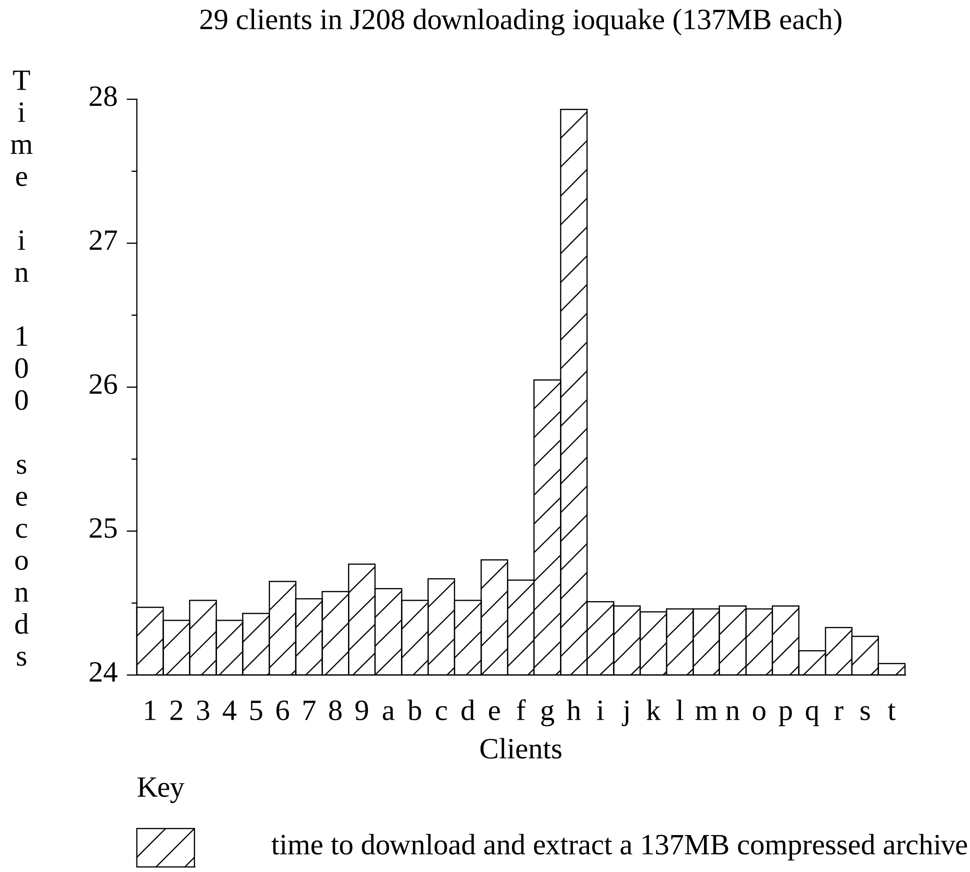

Figure 4: performance of 29 clients downloading and extracting 137MB archive

The archive extraction test cases all involve the clients downloading the archive (from `floppsie.comp.glam.ac.uk`). The data is held on the NFS mounted `/home` directory. So in effect the data is copied back to the `mcgreg.comp.glam.ac.uk` fileserver. The client continues by decompressing and extracting the archive (the data will approximately expand by a factor of 2). Again all data is stored on `mcgreg.comp.glam.ac.uk`. This activity generates a huge bottleneck on this server as all clients try and write data simultaneously.

Conclusion

In conclusion it is encouraging that there were no significant cases of network starvation in J208.

Perhaps the more interesting experimental data is the initial test which shows that the clients can download around 10MB/sec from the local servers. Given the reported Windows logon delay of 40 minutes - this amount of time would allow a client to, in theory, download 24GB of patches. Thus we can conclude that if the actual Windows patch updates are less than 24GB then the proposed cable and switch upgrade is unlikely to deliver a significant performance benefit. One might ponder on whether having a domain server on our local area network might improve (reduce) logon times.

However this is conjecture and ideally it is desirable to conduct further measurements. An interesting experiment would be to substitute a router connecting the lab machines in J208 to the backbone with a GNU/Linux machine running a packet analysis tool `netwatch`. The purpose of this experiment to discover the actual amount of data coming from the Windows domain server and the throughput of this data through the router. Clearly if this is less than the capacity of our departmental network then upgrading the local area network infrastructure will have little effect on logon times. It is also worth remembering

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''. The 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.

[1] Andrew Tanenbaum, Computer Networks, 3rd edition, p.564