

## 6. Bandwidth

As it is an expensive resource, organizations should ensure that they are purchasing the right amount of bandwidth, and test whether they are getting what they pay for. A university should also aim to use the available bandwidth as best it can, not only by limiting abuse but also by providing access to as many people as possible in order to maximize the use of the resource. For example, if there are periods (such as during weekends) when there is low usage of bandwidth, a university could consider providing free or paid access to the local community. Unused bandwidth is wasted money.

### 6.1 Bandwidth testing

Bandwidth testing is an issue for both management and accountants, because when institutions buy an expensive commodity like bandwidth they need to determine if they are getting what they are paying for. ISPs might neglect to explain properly the difference between Committed Information Rate (CIR) and Burst Excess (BE). In particular, the link might be sold on the BE capacity of the link, without the organization realizing that it is not getting that amount of bandwidth some (or even most) of the time. This happens because the ISP's other customers are contending for the same bandwidth.

In a shared system, CIR is the guaranteed amount of bandwidth that the institution should get, while BE is everything above that, depending on how much traffic the ISP's other customers are generating. A shared system is generally better value for money, because more bandwidth than if only CIR was obtained would be available some of the time. However, some CIR is necessary in order to guarantee some bandwidth all of the time. It is also necessary to determine whether the institution is getting its fair share of BE, because if there is too much competition for the BE, it is not worth having a shared system.

Typical link speeds\* (in order of speed):

Ethernet Local Area Network	100 Mbps
T3, DS3 (North America)	44.736 Mbps
Older type Local Area Network	10 Mbps
T1, DS1 (North America)	1.544 Mbps
DSL (ADSL)	512 Kbps to 1.544 Mbps
Typical African university's Internet connection (in 2003)	512 Kbps to 1 Mbps
ISDN	64 to 128 Kbps
Dial-up modems	Up to 56 Kbps

There are tests that can be used to give users or librarians an idea of how much bandwidth they are getting – for example, the Web-based test at <http://www.bandwidthspeedtest.com>, or MicroTik's free Windows-based bandwidth testing utility (see <http://www.mikrotik.com/download.html>). If there is congestion on a link (i.e. if other people are using it) the user cannot expect to achieve the full bandwidth that the institution has access to. Therefore, accurate tests can really be done only overnight. These tests are described in Appendix A.

### 6.2 How much bandwidth is needed

Institutions will need to increase their bandwidth from time to time. Typical reasons include:

- Student numbers tend to grow, and universities increase the number of computers they own.
- The volume of resources on the Internet keeps growing, and tends to become ever more bandwidth-hungry.

\* See Glossary for an explanation of the difference between data rate and link speed.

- New services on the Internet, such as streaming media, may present new opportunities for education, though it is also possible that streaming media will prove to be useful for entertainment only.

The price of bandwidth is falling, even in developing countries where access is only possible via satellite, so increasing bandwidth is more achievable. However, few institutions undertake studies on how much bandwidth is needed. Normally, an institution simply gets as much as it can afford, or increases its bandwidth because a new cheaper rate enables them to do so, or because Internet access simply becomes too slow. There is nothing wrong with this approach, but clear thinking on the subject enables authorities to decide whether they need more bandwidth or whether they need more control over usage (or both). What follows are some pointers from Hughes Network Systems, the MIMCOM network, and Moratuwa University.

A graph of bandwidth required (and bandwidth required per user) against number of users can be expected to look something like this:



Two reasons why bandwidth per user reduces (but overall bandwidth demand increases) are as follows:

- The higher the available bandwidth, the less is needed per user because requests can be satisfied faster, leaving the 'big pipe to the Internet' open for the next user. Also, the waste of bandwidth due to retransmissions is reduced.
- Not all connected users use bandwidth all the time, because in between requesting Web pages they are also reading them.

For these reasons, an ISP such as Hughes Network Systems make the following assumptions for their DirecPC satellite Internet service:

- **175 subscribers per 128 Kbps**, which gives 0.73 Kbps per user. This is the same as the MIMCOM network, which currently has between 1000 and 1200 users. The Hughes network can make use of such a low bandwidth per user because it has so many users (see the graphic above).
- **10% of subscribers will be logged on, and 5% of logged-on users will be active – actually using bandwidth.** This is not a good assumption for the MIMCOM network, where many more than 5% of users are active during peak times, and some sites make large data transfers.

An amount of 0.73 Kbps per user is not much, and in the MIMCOM network it was felt that this should be increased to about 1 Kbps per user; but the MIMCOM network has only about 1400 users. In a larger network, 0.8 Kbps may well be enough if sufficient control and optimization is in place. Of course, the amount of usage also depends on the kind of activity undertaken.


In Moratuwa university, it was found that 1 Kbps per active user (i.e. someone who is online and using the Internet) is too slow, while 3 Kbps per active user is generally sufficient (this can be increased outside working hours).

If users can be limited to a certain amount of bandwidth, deciding how much bandwidth is needed can be done simply by multiplying the number of active users during peak time with the bandwidth that was deemed sufficient. Therefore, if there are 1000 active users during peak time, the university would need 3000 Kbps.

The figure of 3 Kbps was arrived at by Moratuwa through experimentation and deciding what is acceptable.

**6.2.1 Avoid ‘flat topping’**

The graphic below shows a highly congested link at Malawi College of Medicine. The college has a 128 Kbps link to the Internet, and the peak rate is 128 Kbps throughout the working day. When the graph ‘flat tops’ like this, it means that some users have to wait a very long time for anything to happen, and a bandwidth upgrade is needed (unless the congestion is being caused by unnecessary traffic). It also means that bandwidth is being wasted, as applications such as Web browsers have to keep retrying to do the same thing.



Keeping graphs like this one (or the raw data that produced it) will enable the network administrator to know how congested the connection is and how it changes over time.

**6.2.2 Be within the CIR 80% of the time**

In the previous example, the College had a 128 Kbps link to the Internet, without the possibility of BE. In a shared network, it works a bit differently.

In a shared network, it is a good idea to be within the CIR 80% of the time (during the working day). For example, if a network has a CIR of 768 Kbps and a BE of 1 Mbps (and if the network team has properly dealt with abuse and unnecessary traffic), the network could use around 768 Kbps 80% of the time, and up to 1 Mbps for the remaining 20% of the time. This would be a network that is making good use of the available bandwidth, without ‘flat topping’. In the example below, the usage seldom exceeds 100%; when it does, it means that the network is using BE. This network has enough bandwidth, because it is using BE less than 20% of the working day.

