

Introduction

Don Preuss,¹ Dmitriy Beloslyudtsev,² and Steve Sherry³

Created: December 21, 2010; Updated: January 4, 2011.

As the scale of data in next generation sequencing has grown, there is a need to understand the engineering and components involved. In downloading data from NCBI, we have found that many sites are not aware of the numerous variables involved in successfully downloading large quantities of data. To assist in having a more successful download process, this best practices guide has been created.

The quantity of data that can be downloaded from NCBI has increased by orders of magnitude. As such, what was once a fairly casual download with ftp or http, now requires significantly more engineering and attention to detail at all points along the way. While NCBI can assist in some of these, many of the issues you will face are “last mile” issues. That is, issues from your organization’s boundary till the point where data is resting on your storage system. The data transfer will only travel as fast as the slowest link, and while this is simple to declare, often times finding that link can be time consuming. The involvement of your network, security, storage and server teams will often be needed, both in your local organization and part of the central IT group.

In downloading the SRA data, file sizes tend toward the very large, that is 30GB – 1TB. This requires engineering the connection crucial to accomplish downloads in reasonable timeframes. With good engineering, rates of 300Mb/sec are common and speeds of 2.5Gbps can be achieved. This can be two orders of magnitude better than is seen with ftp data transfers.

As a guide point, a 1Gb connection can optimally provide downloads of 9TB per day. Real world observations are lower, so with a 1Gb connection you are likely to transfer around 4TB per day.

Quick Start

In order to get best performance, there are a few common problems that should be checked. If these don’t provide sufficient speed, please proceed further down to the detailed version.

1. If you are using the Aspera client, make sure it is the current version.
2. Ensure all your connections and storage are fast enough, capable of passing 500Mbps (or 50 Megabytes/sec). The likeliest culprits for not achieving these rates are the storage subsystem and the local and building network connections.
3. Your organization’s internet connection must provide sufficient bandwidth. For a business, 45Mbps is not uncommon, so that will be a limiting factor. Many universities have higher connection rates, from 100Mbps to 1Gbps. You should have at least a 100Mb connections, but for transferring large amounts of data, a 1Gb connection is necessary and a 10Gb connection to a research and education network is preferred.

4. Check firewalls for performance. Older firewalls may support physical connections at 1Gb, but may not pass traffic at expected rates. You will need to do a little research to find the limiting rate. Additionally, for Aspera use, certain ports need to be opened in the firewall so that traffic can pass (in the 33001-33010 range for UDP.)

More detailed version

As in the short version above, you want to make sure all your connections and storage are fast enough, capable of passing 500Mb/sec (or 50 Megabytes/sec). The likeliest locations you will find problems are the storage subsystem and the local and building network connections. With Storage, one of the first things to check is the performance of access to the storage devices. Network access to storage is often one of the first limitations. NFS tends to top out around 300MB per second and Windows at around the same rate, but both are often found to be an order of magnitude less. For linux, tuning can drastically improve the performance, for Windows, an upgrade to Win7 (or Windows 2008) is valuable.

For a business, a 45Mbit internet is not uncommon, so that will be your limiting factor. Many universities have higher connection rates, from 150Mb to 1Gb. That will be your limiting factor.

Older firewalls may connect at 1Gb, but may not pass traffic at the higher rates. An upgrade to a recent firewall (which does not have to be that expensive) can quickly provide the additional bandwidth. Note: for Aspera's use, certain ports may need to be opened (in the 33001-33010 range for UDP)

The following list (See Figure 1) graphically shows the common points which have an effect on data movement. We will then discuss each point and provide suggestions for testing and improvement.

NCBI connection points as a reference (these are continually improved to provide good performance):

1. NCBI Storage – High Performance Cluster Storage with multiple 10Gb connections to the NCBI backbone
2. NCBI Network – 10gb connectivity or better to storage and key servers, and 20-40Gb to the core
3. Upload/Download Server – 10gb card with multi-lane PCI and a current multi core processor. Optimized OS settings to support high bandwidth. For linux, see Red Hat tuning guide(1), for Windows the 2008 Tuning Guide(2)
4. Server to Network – 10gb connection to switch
5. NCBI Border – Current technology with 2 – 10gb connections
6. Firewall – Redundant firewalls with 10gb connections
7. Internet II connection – 10gb, Internet I – 1gb(3)
8. If you have an internet 1 connection, you will be limited in download performance as we currently have a 1 Gb connection – contact us to discuss Internet II options

Outside of NCBI:

1. In order to support the increasing speeds for high performance data transfer, it is recommended to have a link to a research and education network rather than the commercial internet. In the US, that would be Internet II. Internationally, there are many links to the US via Internet II, each link generally supporting at least 2.5Gb/sec. You can check (map reference) for access to your country
2. The connection from Internet II to your Organization should be 1Gb or better
3. Network -> External Router – In general, from the entry point of your organization all the way through to your storage, you should support a minimum of 1Gb network connections.

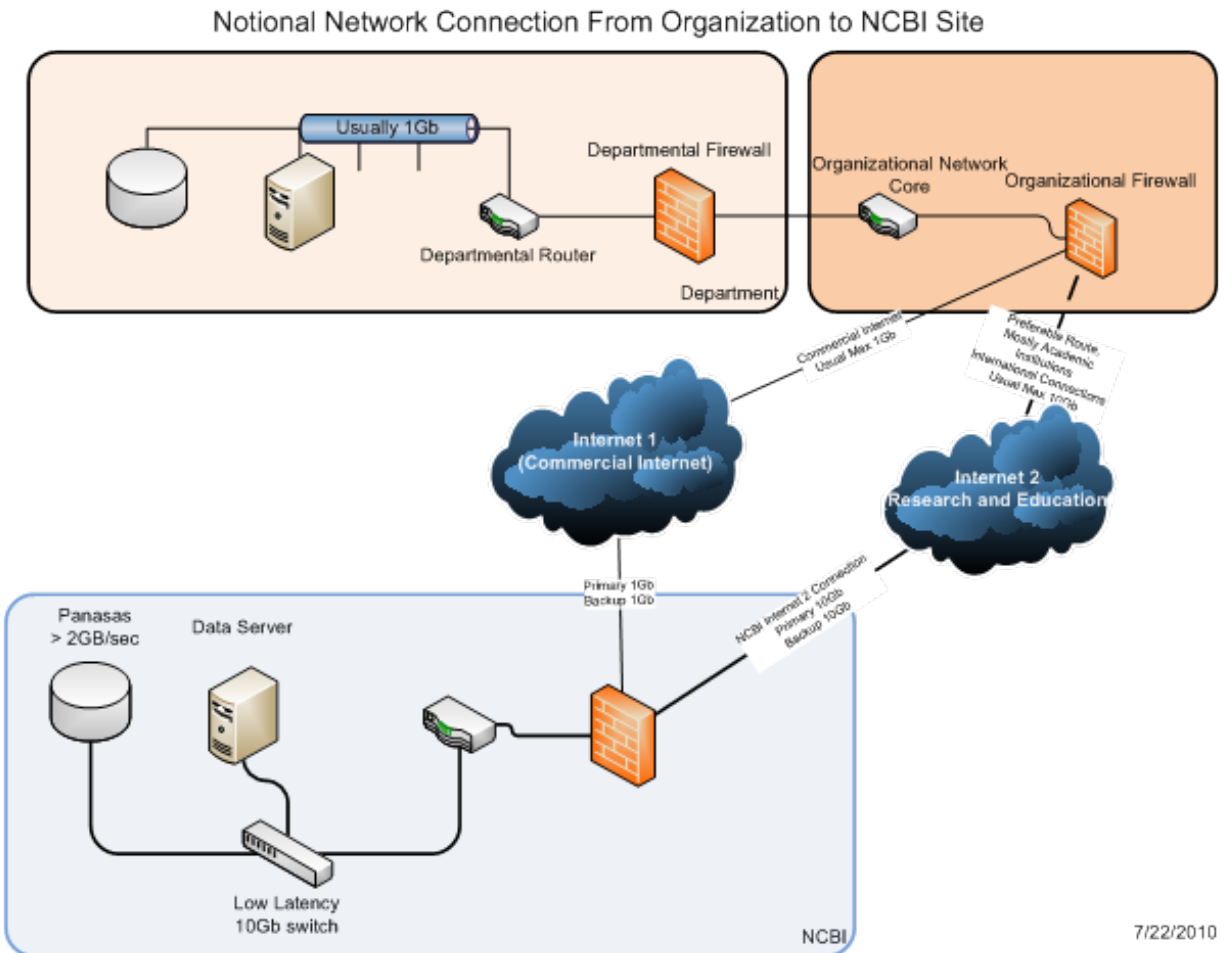


Figure 1: Notional Network Connection an Organization to NCBI

4. External Router – Must be able of supporting the data transfer speeds. It should be noted that a router that supports 10gb interfaces may only have 1-2Gb throughput. This is important and is worth checking.
5. Router -> Firewall – must be at least as fast as external connection
6. Firewall – Must be current technology, support UDP data transfer at external speeds (greater than 400Mb). This has often been a bottleneck at other sites. Just as with the router, firewalls that supports 10gb interfaces may only have 1-2Gb throughput. This is important and is worth checking (See Figure 2).
 - a. Ensure the UDP ports 33001-33010 are permitted, that the timeout (if there exists one is longer than a data transfer, say 48 hours)
 - b. If NATing, make sure the NAT timeout is longer than 48 hours.
 - c. Ensure UDP throttling is not enabled. This will cause a connection to terminate after a set period of time or set amount of data transferred.
7. Central Network -> Department – This network connection could be slow, sometimes 100Mb. It must be at least 1G, preferably 10Gb. Some organizations limit bandwidth at this connection, so even though the physical connection may be 1Gb, there may be policies that limit access. You should check with your organization's network staff. On a campus, there may be value in a direct connection to the Internet II border firewall rather than upgrading the entire infrastructure on campus which might not require higher bandwidth.

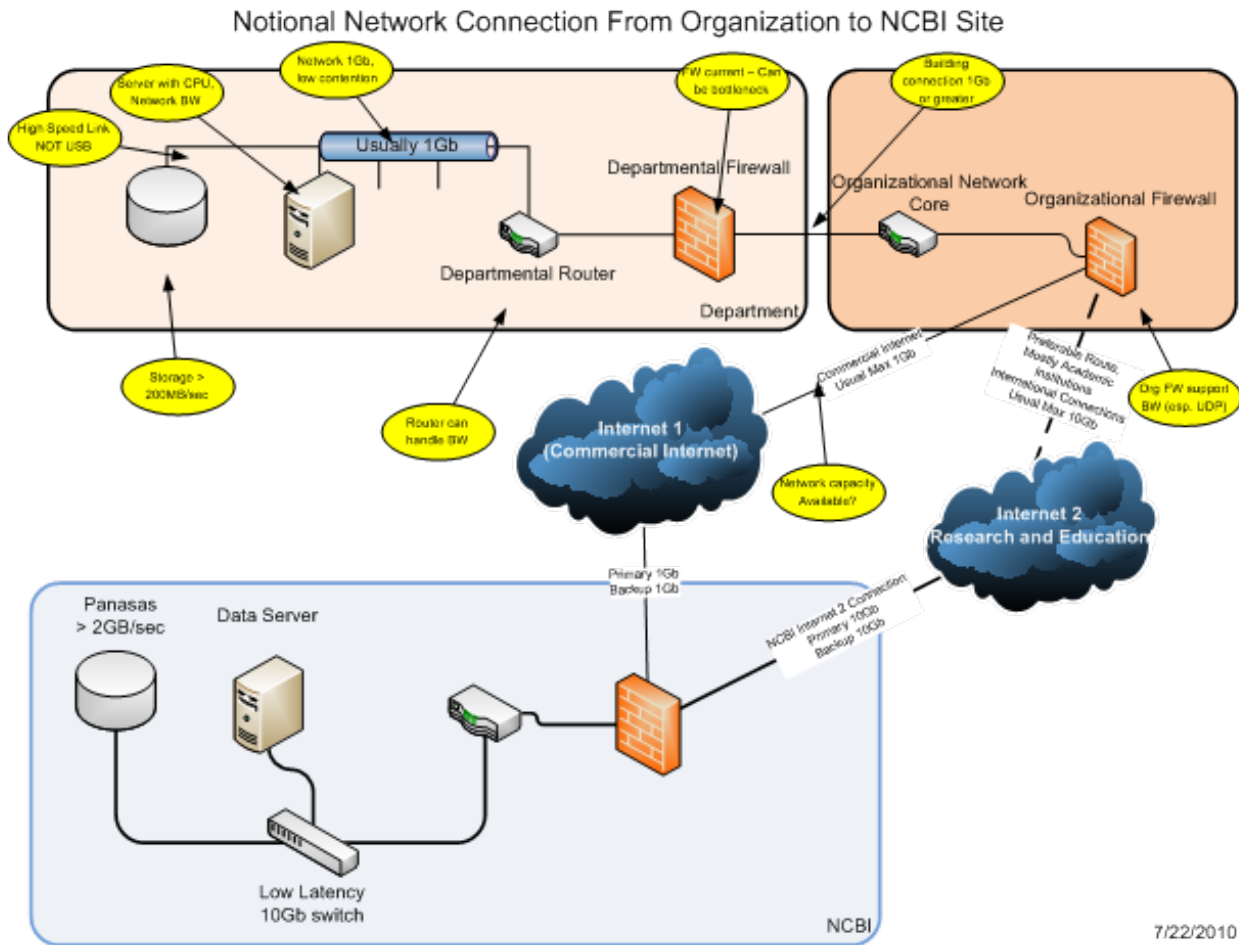
8. Departmental Firewall – This is often a slow point, one must ensure performance matches connection speeds. As above, be aware that some firewalls support a port speed higher than the actual processing ability. (example: Cisco ASA has 10Gb ports, but only supports 5Gb/sec with standard packet sizes)
9. In building network – should be 1Gb or better connecting the firewall, servers and storage.
10. Servers should be connected with 1Gb interfaces or better, and one need ensure there is no overallocation on the network switch. Further, even though monitoring of the network interface may show additional capacity exists, migrating to a higher performance 10gb connection has been shown to double performance on local networks.
 - a The network configuration defaults for linux need to be increased for high performance. Please see the redhat tuning, PSC tuning guide or DoE Fasternet.es.net/tuning.html.

```
# increase TCP max buffer size
net.core.rmem_max = 4194304
net.core.wmem_max = 4194304
# increase Linux autotuning TCP buffer limits
net.ipv4.tcp_rmem = 4096 87380 4194304
net.ipv4.tcp_wmem = 4096 65536 4194304
```

- 11 Storage subsystem – must be able to support 200MB/sec. with an internal RAID array or an external, high performance storage system (NAS). **Note: Externally connected USB drives will often provide very poor performance and are not recommended.** The storage subsystem is often the critical step in having a high performance download.
 - a On Linux, for local storage, alternative filesystems may provide better performance. Better performance has been seen with xfs, especially when there are large numbers of simultaneous writes and reads.

References

1. Redhat Enterprise Performance Tuning Whitepaper
<http://www.redhat.com/docs/wp/performance-tuning/iotuning/index.html>
2. Windows Tuning Guidelines for Windows 2008
http://www.microsoft.com/whdc/system/sysperf/perf_tun_srv.msp
3. <http://www.internet2.edu/>
<http://www.psc.edu/networking/projects/tcptune/>



7/22/2010

Figure 2: Notional Network Highlighting Common Performance Bottlenecks