THE FEASIBILITY OF MOVING TERABYTE FILES BETWEEN CAMPUS AND CLOUD

Adam H. Villa
Department of Computer Science
University of New Hampshire
Durham, NH, USA
email: ahvilla@cs.unh.edu

Elizabeth Varki
Department of Computer Science
University of New Hampshire
Durham, NH, USA
email: varki@cs.unh.edu

ABSTRACT

Cloud/Grid computing is envisioned to be a predominant computing model of the future. The movement of files between cloud and client is intrinsic to this model. With the creation of ever expanding data sets, the sizes of files have increased dramatically. Consequently, terabyte file transfers are expected to be the "next big" Internet application. This application is different from other Internet applications in that it requires extensive bandwidth, orders of magnitude larger than the bandwidth requirements of existing applications. It is essential to determine whether or not existing network infrastructures can handle the augmented workload that terabyte transfers would create. This is particularly critical for academic campus networks that are already under strain from high user demand. The paper evaluates the system level challenges of incorporating terabyte transfers into an existing campus network. The evaluation finds that large file transfers can be handled by the current campus network without making major changes to the infrastructure. It is vital to employ a system level service that schedules and monitors the terabyte transfers on users' behalf. By removing control from users, the service is able to leverage low demand periods and dynamically repurpose unused bandwidth.

KEY WORDS

Large file transfers, data retrieval, clouds, academic networks

1 Introduction

A multitude of scientific and commercial applications routinely generate terabyte and petabyte sized files. For example, the Large Hadron Collider at CERN in Geneva generates copious numbers of large files that researchers around the world are eager to access [1, 2, 3]. Due to the enormity of these file sizes, transferring them over shared networks is not trivial. In order to electronically transmit files from CERN's data grid to an academic institution on another continent, the data have to travel from grid storage to the researcher's storage system over the Internet. Starting from grid storage, the desired file travels over the private grid network and over public networks until it reaches the destination campus network. From there the file will finally arrive at the researcher's local storage. The path of the retrieved file may include several networks, which are controlled by various organizations. This paper analyzes the issues and challenges involved in moving very large files over a campus network that is shared by thousands of concurrent users running a variety of applications. A campus network, with its users and applications, is a microcosm of the Internet, so we expect the issues to be relevant to other shared network environments.

Currently, only a small percentage of academic users, mainly researchers in disciplines like physics and biochemistry, need access to large data sets stored elsewhere. Since electronic transmission of large data sets is difficult, these researchers often transfer their data via hard disks transported by snail mail [4]. In rare cases, fast links can be manually and temporarily set up between two locations for transfers of large data sets by working with network administrators. As large scientific and commercial data sets become available in a growing number of disciplines, a greater number of academic users will require access to these data files. Since these data sets are often multiple petabytes in size, researchers will often require subsets of the data with file sizes in the range of hundreds of gigabytes to several terabytes. Users routinely require and prefer local access to these files for processing and other tasks. Even the output of remote computations in both grid and cloud environments can be in the same magnitude of file size. The movement of large private files between the cloud/grid and its clients is a commonplace occurrence. Therefore, in addition to specific research groups, individual users on campuses will require access to large files. Current trends in computing and the increasing sizes of files predicate the need for efficient techniques to transfer large files to and from campuses.

Large file transfers impose a much higher bandwidth burden than any other application. Users want to be able to retrieve/transmit large files quickly with a click of a mouse, without having to worry about errors and retransmissions. Software tools for file transfer include FTP, GridFTP [5], Globus File Transfer [6, 7], HTTP file transfer, BitTorrent [8] and email. It takes some effort, more than a mouse click, to use these tools for larger file transfers, especially in heavily utilized campus networks [9, 10]. In order to move files quickly, the campus user must have fast links. For ex-

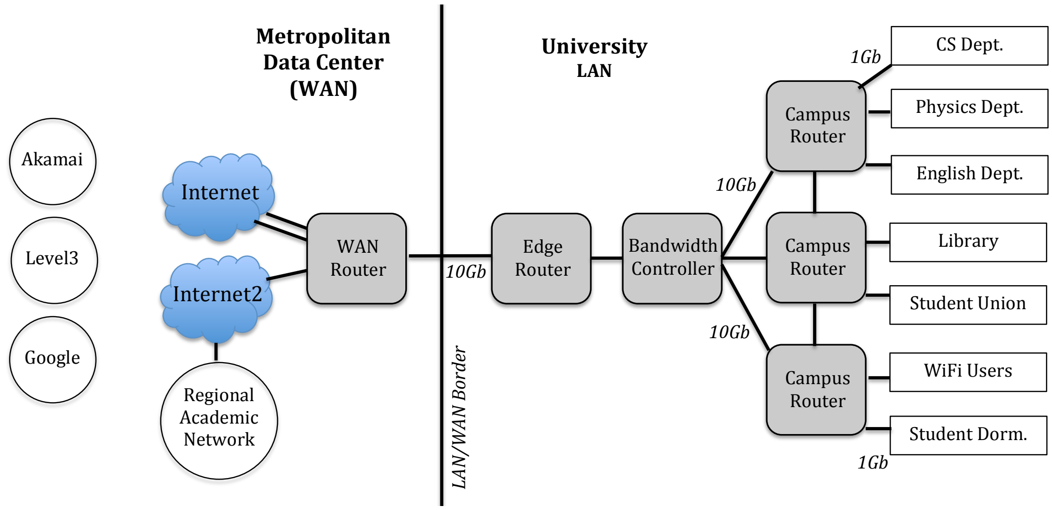

Figure 1. Network layout for the campus network and its connection to the shared data center in a nearby metropolitan area.

ample, to transmit a 1 terabyte file over a 1 Gb link would take at least 2.3 hours, and over a 5 Mb link would take at least 19.4 days. However, increasing the bandwidth for large file transfer users would limit the available bandwidth for other users. Satisfying the performance requirements of large file transfers is important, but it should not inhibit the performance of other applications.

There is considerable research interest in techniques for large file transmission [11, 12, 13, 14, 15]. The majority of existing research focuses on new network hardware and new network protocols for large file transfers [16, 17, 18, 19, 20]. Purchasing new hardware for a particular application may not be cost effective. Before investing in new hardware for large file transfers, it is prudent to investigate whether the existing campus computer and network infrastructure can be parlayed to support large file transfers. An efficient solution must not only ensure that the performance requirements of users transmitting large files are satisfied but also that the addition of large file transfers does not impede other users and applications. The problem of adding a high-load application such as large file transfers to a shared network environment is not just a network issue. It is a systems issue and requires an understanding of the milieu in which these transfers occur. The users transmitting large files share the campus network with myriad users running a variety of applications with different performance requirements. In order to satisfy the performance requirements of all network users/applications, it is necessary to understand the issues and challenges of incorporating large file transfers into the existing campus design. The contribution of this paper is the systems level evaluation of both the potential and the limitations of campus Internet infrastructure for large file movement.

The focus of this paper is not the development of a new tool for large file transfer. Instead, the focus is on the feasibility of using existing tools to transfer large files to and from overloaded campus networks. Our evaluation shows that large file transfers place a heavy load on the network that may translate to sub-par performance of other applications. Therefore, large file transfers must be monitored, and these transmissions should not be controlled by end users. We conclude that campus networks can handle large file transmissions if control is placed in the purview of a systems agent, such as the border controller, which has a bird's eye view of the network traffic at all times and can take advantage of low load periods and dynamically utilize unused bandwidth.

2 Campus network

Moving large amounts of data across the campus network will place a heavy burden on system resources. It is therefore necessary to examine the infrastructure of the system in order to determine the feasibility of accommodating large file transfers for multiple users.

We present the infrastructure and configuration of the campus network at our midsize university. This academic system is a microcosm of the Internet with over 10,000 users including students, faculty, researchers and staff. All of these users are utilizing a variety of applications and transferring different types of data, similar to the Internet as a whole. We recognize that our network only represents one example of a campus network, however many other universities have similar setups and configurations. All campus networks have a local area network that utilizes shared connection(s) to a greater wide area network and the Internet. It is the interaction between the LAN and WAN that is critical, as this is often the bottleneck for many networks [21, 22].

Our campus network supports both Ethernet and WiFi connections for thousands of users in all buildings across campus. These connections are grouped into multiple subnets throughout the local area network as illustrated in Figure 1. Each of these subnets is connected to the greater campus network via 1 Gb links. Campus routers receive data from the individual subnets and any external traffic is forwarded to the edge routers of the campus network. The core of the campus network is connected with 10 Gb links.

Traffic destined for external locations outside of the campus network must pass through the border between the LAN and WAN. This border is managed by a bandwidth controller device. This device monitors and adjusts data streams passing through the border. After the aggregated external traffic passes through the border, it makes it ways towards the WAN through a private 10 Gb fiber link to a nearby metropolitan area. When it reaches the city, the traffic arrives at a shared data center, which contains access points to major telecom networks, regional universities and major corporations' services (such as Google, Akamai, Level3, etc.). The outgoing data are then routed to three different wide area network connections. Two of these connections are to the public Internet and one of these connections is to Internet2, a non-profit network designed to support research and educational institutions. Traffic destined for the Internet is load balanced between the two general Internet WAN connections, which have a total bandwidth capacity of 1.5 Gb/s. The Internet2 connection has a variable bandwidth capacity, which allows on average 500 Mb/s. This results in a total bandwidth capacity of 2.0 Gb/s for the entire university network. The private connection to the data center has the ability to support up to 10 Gb/s, which allows for future expansion and specialized projects.

Any incoming data that is destined for a user on the campus network arrives at one of the three WAN connections. This data could be streaming multimedia from a nearby CDN server or a file server at a regional university. All of this data crosses the LAN/WAN border and passes through the bandwidth controller before being routed to the correct subnet and finally to the end user.

At our university, there are over 10,000 registered users and the network is normally handling 6000 concurrent users actively using the WAN connections. In order to guarantee that each of these 6000 users will be able to have fair access to the shared Internet connections, the university employs a bandwidth controller. This device sits at the border between the LAN and WAN. The bandwidth controller ensures that a user only has access to a portion of the available bandwidth. Campus users are restricted to a maximum bandwidth of 8 Mb/s. As more users utilize the shared Internet connections and demand for bandwidth increases, the per user bandwidth allowance will be further restricted to ensure that each active user has an equal portion.

The bandwidth controller is crucial to ensure that everyone has access to the shared Internet connections. It however does not guarantee that users will have sufficient bandwidth and capabilities to utilize their desired applications. Under high load conditions, a user might only receive a small fraction of available bandwidth. This amount might suffice for web browsing and email messaging, however streaming media and applications requiring low latency or quick response times will suffer. Applications, such as large file transfers, will have low throughput and require significant time to complete data transfers. At the maximum allowed rate of 8 Mb/s, it would take a user over 12 days to retrieve a 1 TB file at that constant rate. During peak times this rate will be even lower, which will extend the duration of the transfer.

Feasibility: In order to accommodate large file transfers, system resources need to be able to handle the increased burden created by these workloads. After examining the infrastructure and configuration of our campus network, which has a similar structure to other campus networks, we find that it is capable of supporting large file movements without significant modifications to the infrastructure. The core of the campus network and the link to the WAN connections is 10 Gb, which would theoretically allow the transfer of a terabyte file in under 15 minutes. The links to the end user on campus could support a maximum of 1 Gb/s, which provides a theoretical time of 2.3 hours. If all of the connections in the data path are able to support these rates, then the storage systems will be the limiting factor of the transfer rate. The border controller that manages the interface between the LAN and WAN has a bird's eye view of all traffic passing through the border. Since it has complete knowledge of the workload present in the system, it could be utilized to schedule large file transfers and to allow these tasks increased bandwidth in order to complete quickly.

3 Traffic on the campus network

In the previous section, we examine the existing infrastructure and configuration of the campus network. We identify the resources that would be available for large file transfers. In this section, we examine the utilization levels of these resources to see if they can accommodate additional workloads from large file transfers.

We recognize that the traffic we observe represents only our specific campus network and its users. Obtaining the following detailed data about bandwidth usage and user workloads required several rounds of authorization and working with network administrators to access live, mission critical hardware devices. Attempting to obtain similar in-depth data from other colleges and universities proved impossible due to security concerns and confidentiality issues. We realize that the specifics from our analysis might only relate to our network, but the trends that we observe are definitely present at universities throughout the country [23, 24, 25].

In order to monitor the usage on our network, we utilize network monitoring software and hardware devices that examine the raw data passing through the LAN/WAN border and the bandwidth controller. We aggregate all of this information together to get a comprehensive view of the traffic coming in and out of our network. Using this in-

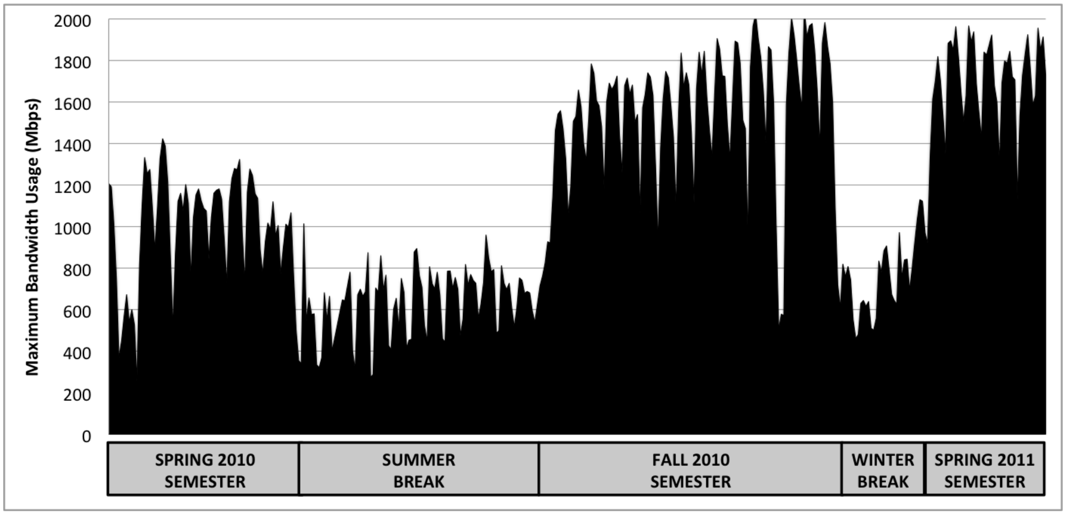

Figure 2. Changes in maximum bandwidth consumption for the past 12 months for all data passing through all of the university's shared Internet connections. Each semester user demand and bandwidth consumption increases.

formation, we conduct a thorough analysis and determine several key findings regarding the demand and user workload placed on our campus network.

Bandwidth usage: We examine the changes in campus bandwidth consumption over the course of several semesters. Figure 2 demonstrates the variations in the daily maximum bandwidth consumption for multiple semesters. When classes are in session during fall and spring semesters, there is significant demand for bandwidth on campus. Regardless of class schedules, there is a minimum level of demand continually present as the university hosts several governmental projects that routinely transmit and receive replicated data sources.

The major consumers of bandwidth on campus are student users. The enforced total bandwidth limitation for these users is 1.5 Gb/s and all data for the entire campus network passing through the shared WAN connections are currently restricted to 2 Gb/s. The network has the potential capacity for 10 Gb/s, but the WAN connections are presently not configured for these rates. Figure 2 illustrates that campus bandwidth usage frequently passes the 1.5 Gb/s mark and actually reaches the top limit of 2 Gb/s several times during the Fall 2010 semester, as seen as peaks in the graph. Even with these occasional high levels of demand, there were still periods of low utilization, as illustrated by the valleys in the graph. During these low demand times, available bandwidth could be devoted to large file transfers.

In addition to variations in bandwidth from semester to semester, we also observe changes in bandwidth consumption depending on the day of the week. Figure 3 shows the bandwidth usage (maximum, average and minimum) for a typical week during the current Spring 2011 semester. We identify that network usage is at its highest between Sunday evening and Friday afternoon. This correlates with classes beginning and ending for a typical week. The time period between Friday night and Sunday afternoon has the lowest levels of network utilization. During this period, even the maximum bandwidth rates are significantly decreased in comparison to a weekday. We attribute this occurrence to the fact that many students and staff leave campus or reduce their network usage on the weekends. These low demand periods leave available bandwidth that could be utilized for large file transfers.

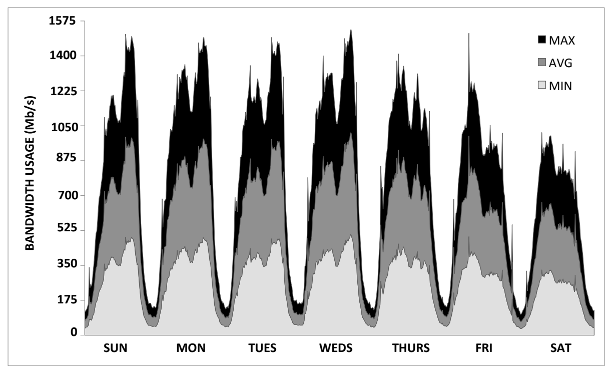

Figure 3. Changes in the minimum, average and maximum bandwidth usage (all receiving and transmitting traffic) for a typical week during the Spring 2011 semester.

We further examine the changes in bandwidth consumption from an hourly level. We observe usage patterns based on the time of day. In Figure 4, we examine the bandwidth usage (maximum, average and minimum) for each hour in a typical day during the current Spring 2011 semester. Peak bandwidth consumption transpires between noon and midnight. There is a decrease in demand around dinnertime and then consumption increases until

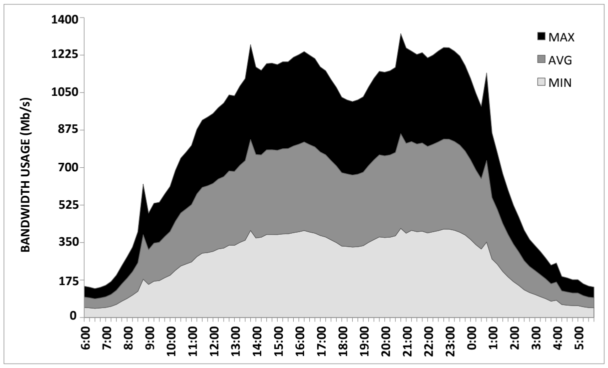

Figure 4. Changes in the minimum, average and maximum bandwidth usage (all receiving and transmitting traffic) for a typical day during the Spring 2011 semester.

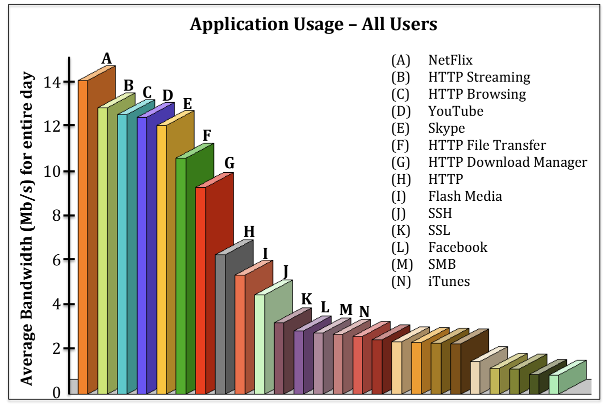

Figure 5. Most active protocols utilized on a typical day for all users.

about 1AM when demand then starts to decline. Between 4 and 7 AM, we observe the lowest amounts of bandwidth consumption for all users. After 7AM, demand begins to rise steadily as faculty/staff users return to campus and students prepare for their day. We identify from the minimum bandwidth consumption values that there is always demand on our campus network regardless of the time of day, however there are most definitely periods of the day when idle bandwidth is available. This unused resource could be repurposed for large file transfers.

Applications on campus: We further examine the workload on the network by analyzing the individual components of the traffic. We monitor the traffic flow for a period of 35 days for all users and extract usage data in order to determine the characteristics of the applications consuming campus bandwidth.

Figure 5 shows the average bandwidth utilized by the top applications/services for all users during a typical day. We find that Netflix and HTTP Streaming video applications consume the most bandwidth on campus for all users. Web browsing, YouTube and Skype follow with the next highest percentages of bandwidth usage. Streaming audio/video and VoIP are by far the largest consumers of resources for all users. These applications are time critical and highly sensitive to increases in network latency. During peak times, when the network is under high load, the addition of large file transfers would not be appropriate since this additional burden would cause congestion and performance problems for these time critical applications.

Unbounded user demand: During our examination, we also find that users bandwidth demands are unbounded and that any available bandwidth given to them is quickly consumed. An example of this situation appears every weekend on campus when the students' bandwidth allotments are increased. Since there is generally lower demand from users, especially from faculty/staff users, network administrators allows students to have larger bandwidth allotments for the weekend. During this time period, there is a significant increase in the bandwidth utilization for all applications when more bandwidth is given to students. Bandwidth consumption for streaming multimedia increases by more than 200% for all users. Skype usage triples and Netflix quadruples in usage. Netflix dynamically adjusts to changing bandwidth conditions and will automatically use additional bandwidth to increase video quality for the user up to high-definition standards. Other universities also experiment with increasing user bandwidth and they have experienced the same complications of bandwidth consumption [23, 24]. Despite the unbounded demand, bandwidth consumption still has periods of low utilization that could be utilized by large file transfers.

Feasibility: We find that users' bandwidth demands are unbounded and users will utilize any bandwidth that is provided to them, especially during peak periods. The composition of user traffic is dominated by time critical applications, such as streaming multimedia, web browsing and VoIP, which are highly sensitive to changes in network latency and congestion. We also find that there are varying levels of demand during different times of the day and on different days of the week. High demand is present during the week when students are actively connected to network, specifically between noon and midnight. When the campus network is under high load, large file transfers should not be placed in the system, as they will negatively impact other applications and will take longer than necessary. Between midnight and noon and on the weekends, the number of connected users is significantly lower and so is bandwidth demand. It is during these low usage periods that large file transfers should take place.

4 Impact of Large File Transfers

In this section, we examine the impact of large file transfers on the existing campus network infrastructure. We identify the issues caused by large file transfers when there are no modifications to the campus network architecture and no additional hardware/software is added to the system.

To demonstrate the performance impacts of moving large files from the cloud to the campus, we run experiments where multiple user machines are used to concurrently retrieve portions of a terabyte dataset. These experiments simulate multiple users retrieving large data files at the same time. Each of the machines used in our experiments is configured to open five parallel data streams to a source server. The source servers are located at nearby research institutions and universities. The data paths from the multiple sources to the campus network utilize all three of the shared WAN connections. Transfers from the nearby universities utilize the Internet2 connection and the remaining sources utilize the two general Internet connections. All of the data for these transfers passes through the bandwidth controller device at the LAN/WAN border, however the device is configured to allow **unrestricted** access to the machines in our experiments. (When these experiments were conducted, the bandwidth controller did not yet restrict research machines in the Computer Science department. This has since been changed and all machines across campus now fall under the control of this device.)

We utilize up to five client machines on a single subnet of the campus network. Since the machines are located on the same subnet, they utilize a shared 1 Gb fiber link to the campus router. The experiments are run on the live campus network and we have no control over the workloads of other users present in the system. Other users' applications are running at the same time and consuming resources. Our experiments are run at all times of the day for an entire week. We average the results of all of our experiments and present those averages.

We find that a single machine using five parallel data streams is able to retrieve data at a rate of 66.5 MB/s on average. When two machines are operating in a parallel, the average total bandwidth consumption for both machines is 85.7 MB/s, an increase of 29%. Adding another client machine results in a combined bandwidth rate of 101 MB/s. When four machines are concurrently retrieving data, the total rate increases to 117 MB/s. This is approaching the maximum limit of the 1 Gb link (128 MB/s). When a fifth machine is added, the combined throughput reaches 123 MB/s.

During our live experiments, we were contacted by the local subnet support team, as well as the university's telecommunication department. Our experiments impacted the service of the subnet as well as the general Internet connections during peak usage times. During these high demand periods, our experiments increased the load of the subnet link to near full capacity, which resulted in service problems for users. During low usage intervals, users' workloads were only minimally impacted due to the limited number of active users on the subnet.

Using bandwidth monitoring devices, we are able to obtain bandwidth utilization graphs for all of the WAN connections. Figures 6a, 6b, and 6c show the bandwidth utilization for the three WAN connections before, during and after our experiments. Before we began our evaluations, the university was on a mid-semester break. During this time, there was very little network load. Only automated and system traffic is present in the system at this time. Once users returned to campus, we initiated our experiments, so that our workload would be intermixed with normal everyday traffic.

Since our experimental traffic is mixed with all other users, we are unable to precisely isolate our traffic in the graphs. Figure 6a illustrates the bandwidth utilization for the Internet2 traffic. Since most student and staff traffic very rarely use the Internet2 link, it is easier to discern our experimental workload in the graph. Multiple sources in our experiments are located at universities on the Internet2 network and therefore our workload has a dramatic impact on the bandwidth utilization of the Internet2 link. As the graph indicates, the normal usage before and after our experiments is quite low in comparison. There is a significant increase in traffic on this link during our experiments.

The impact of our experiments on all WAN connections is slightly harder to see when it is intermixed with all other users' traffic. Figure 6b illustrates the total bandwidth utilization for all of the campus WAN connections. During our experiments, the total utilization reached its highest peaks during the two week time span. When the time range for this graph is increased to four months, as shown in Figure 6c, it is easier to notice the impact of our experiments. The bandwidth utilization again reached its highest peak during our experiments, as well as maintained a higher utilization for the entire experimental period. From all three of these graphs, it is clear that large file transfer workloads can impact the performance of the entire campus network, especially during high demand periods. Even with our experimental workloads, there are still periods of low utilization, as demonstrated by the valleys in the graphs. It is during these low usage periods that we suggest large file transfers should be completed.

Feasibility: From our experiments, we find that it is possible to retrieve large data files over the campus network. We identify that these workloads impact system performance and cause congestion during peak periods. There is no benefit to any user by running these transfers during high load times. Campus users will experience delays and jitter in their time critical applications and the large file transfers will see decreased transfer rates and longer durations. If the transfers are restricted to only operate during low utilization periods however, then the performance impact on user workloads will be minimal and the large file transfers will find faster transfer rates and shorter service times.

5 Potential and limitations

During our system level feasibility study, we identify three key challenges that must be addressed before terabyte transfers can become commonplace on the campus network.

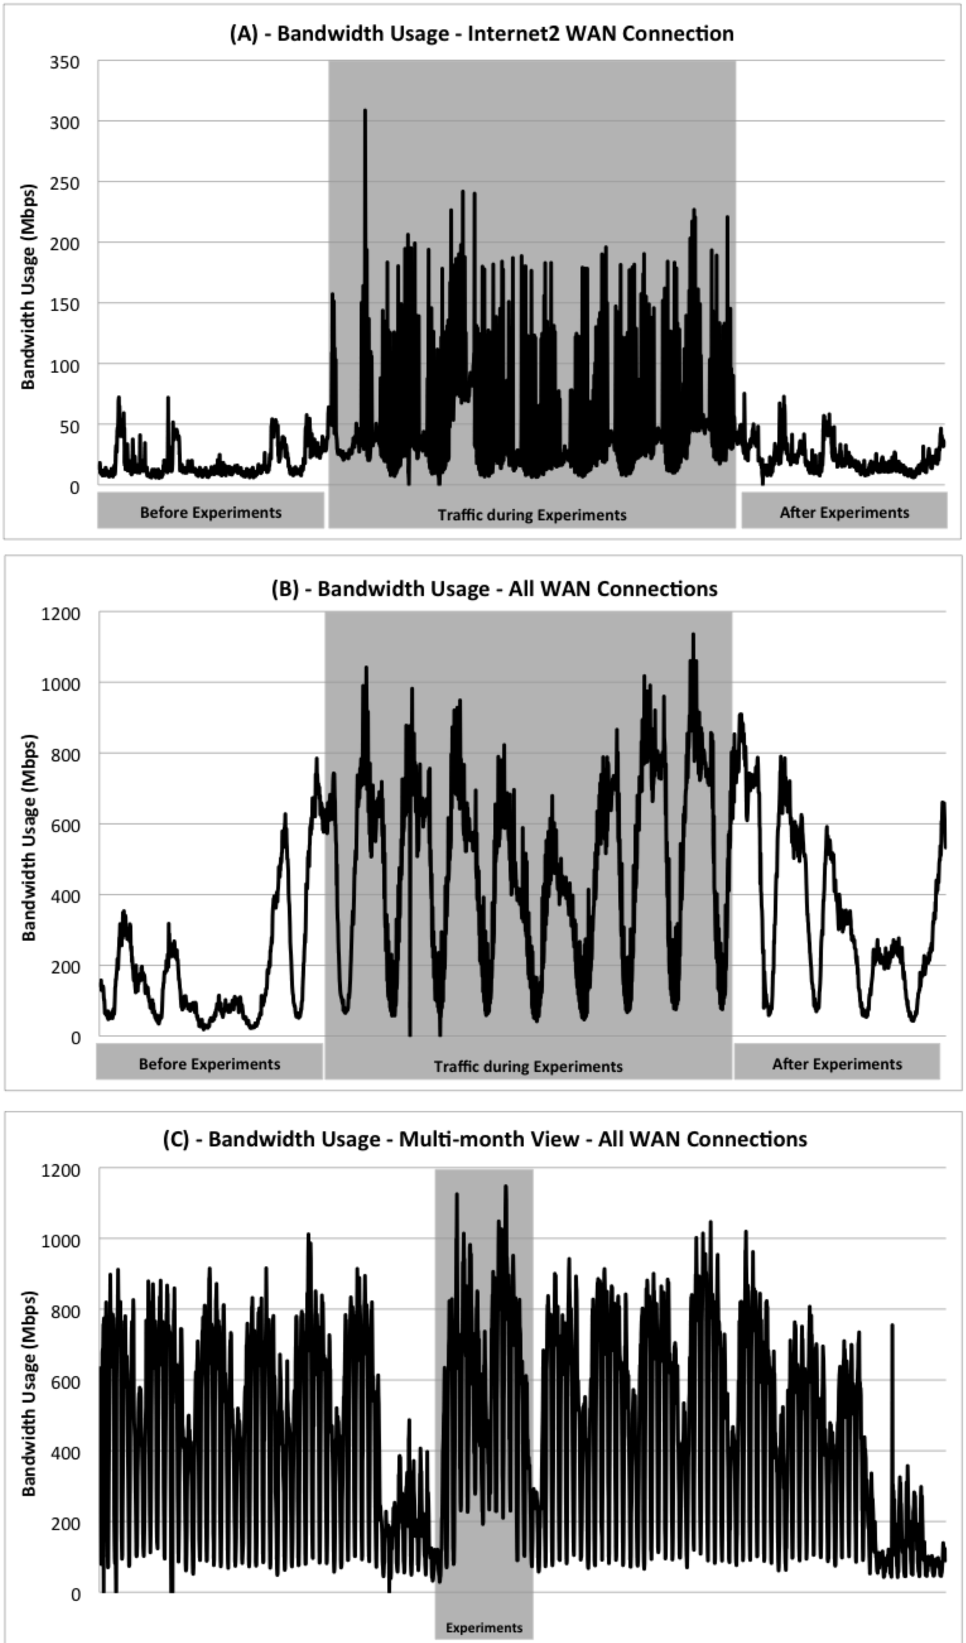

Figure 6. Bandwidth usage on wide area network connections before, during and after our evaluations. The shaded regions indicate the time period during our experiments. (a) - This graph shows bandwidth usage of the Internet2 connection for a two-week period. (b) - This graph shows total bandwidth usage of all WAN connections for a two-week period. (c) - This graph shows total bandwidth usage of all WAN connections for a four-month period.

1. As Internet applications evolve and new services become available, the demand for bandwidth is expected to outpace the available bandwidth on several campuses [23, 24]. In order to ensure fairness for all users, the bandwidth allotted to individual users is limited. Large file transfers have the highest bandwidth requirements in comparison to other Internet applications utilized by users. With these restricted bandwidth allotments, terabyte transfers would take several weeks to complete. On the other hand, allowing unrestricted bandwidth to large file transfers would greatly reduce the bandwidth available for other applications.

2. The majority of campus users are running interactive, time-critical applications. Any loss of bandwidth or congestion can result in jitter and slowdowns for these applications. To a user staring at the "screen," even a small delay can appear endless and frustrating. These services can therefore not be impacted by terabyte transmissions.

3. Users do not have the resources or the computer savvy to handle large file transfers on their own. The current file transfer tools place enormous burden on the end user. New, user-friendly tools designed specifically for terabyte transfers are needed.

We conclude that terabyte transmissions should not be allowed free rein on campuses, but should be controlled by administrative/system software. Our feasibility study also identifies the advantages provided by the campus infrastructure with regard to incorporating terabyte transfers:

- The bandwidth controller placed at the border between the campus LAN and WAN manages all traffic moving in and out of campus. The controller has a complete view of the campus traffic conditions. Moreover, the controller manages the bandwidth given to each user at all times. Therefore, the bandwidth controller has the knowledge and the authority to control the bandwidth given to terabyte transmissions.

- Our studies show that while campus users place heavy load on the network, the load is not consistent during all times of the day. There are periods during each day when there is very low usage of the network. During these times, the network can be specifically employed for terabyte transfers.

- File transmissions are not time critical applications. Users do not want to deal with errors, timeouts and retransmissions, they just want to upload/download files with minimum problems. Therefore, users would be satisfied to have a systems level manager take charge of the transmission process.

6 Conclusion

Our paper studies the feasibility of terabyte transfers to and from campus networks. We conclude that while it is difficult to incorporate this demanding application, it is doable. The solution should ensure that end users do not control when and how these transmissions occur. We recommend that terabyte transmissions borrow from the SMTP (email) protocol - similar to email transmissions being controlled by mail servers, terabyte transmissions should be controlled by "cargo" servers. These servers would be specifically in charge of the large file transfers for the entire campus network and would ensure that transfers have both high reliability and high efficiency. We recommend that a cargo server be placed at the border between the campus LAN and WAN. The cargo server and the bandwidth controller would work cooperatively. When users want to transmit cargo, that is, when users want to transmit large files, they submit their requests to the cargo server that contain information about the files to be exchanged and any other data that is required to complete the transfer. During low usage periods on the campus network, the cargo server would initiate and monitor the users' transfers. The server would also be in charge of sub-dividing the file into smaller pieces if needed for efficient data transfer. Since the cargo server is working in tandem with the bandwidth controller, the controller can allow the cargo server to utilize any unused bandwidth. We envision cargo servers to be placed on campus networks that regularly transfer large amounts of data. These distributed servers could work together in order to efficiently transfer data from users on one network to another. The cargo servers would use medium and low load periods for transfers, so that they can use more of the available bandwidth without impacting other users. Since low demands periods can be quite different for the two ends of the data transfer path, the service allows the transfer to utilize multiple cargo servers along this path to transfer data between servers only when their low demand periods match. Users can check the progress of their requests and would be notified when their requests are complete.

In conclusion, the paper presents the challenges of incorporating terabyte file transfers into the existing campus network infrastructure. We show that large file transmissions should be handled carefully since the load placed by these transmissions can negatively impact the performance of other applications on the shared network. We propose a solution that integrates into the existing infrastructure. The details of our proposed solution are left for future work.

References

[1] L. R. J. Knobloch, "Lhc computing grid - technical design report," CERN, Tech. Rep. LCG-TDR-001, June 2005.

[2] D. Minoli, *A Networking Approach to Grid Computing*. John Wiley and Sons, Inc., 2005.

[3] C. Nicholson and et. al, "Dynamic data replication in lcg 2008," in *UK e-Science All Hands Conference*, Nottingham, September 2006.

[4] "Amazon web services (aws)," *http://aws.amazon.com/importexport/*.

[5] W. Allcock, J. Bresnahan, R. Kettimuthu, and M. Link, "The globus striped gridftp framework and server," in *Supercomputing (SC)*, 2005, p. 54.

[6] "The globus alliance, http://www.globus.org/," 2009. [Online]. Available: http://www.globus.org

[7] I. Foster, "Globus toolkit version 4: Software for service-oriented systems," in *IFIP International Conference on Network and Parallel Computing*, 2006, pp. 2–13.

[8] B. Cohen. Bittorrent, http://www.bittorrent.com. [Online]. Available: http://www.bittorrent.com

[9] J. Bresnahan and et. al, "Globus gridftp: What's new in 2007," in *GridNets*, October 2007.

[10] J. Kurose and K. Ross, *Computer Networking: A top-down approach featuring the Internet*. Addison-Wesley, 2005.

[11] W. Allcock and et. al, "Secure, efficient data transport and replica management for high-performance data-intensive computing," in *IEEE Mass Storage Conference*, 2001.

[12] W. Allcock, "Data management and transfer in high performance computational grid environments," *Parallel Computing Journal*, vol. 28, no. 5, pp. 749–771, May 2002.

[13] A. J. G. Hey and A. E. Trefethen, "The data deluge: An e-science perspective," pp. 809–824, 2003.

[14] M. Li and M. Baker, *The Grid - Core Technologies*. John Wiley and Sons, Inc., 2005.

[15] E. Nygren, R. K. Sitaraman, and J. Sun, "The akamai network: a platform for high-performance internet applications," *SIGOPS Oper. Syst. Rev.*, vol. 44, pp. 2–19, August 2010.

[16] E. P. Jones, L. Li, J. K. Schmidtke, and P. A. Ward, "Practical routing in delay-tolerant networks," *IEEE Transactions on Mobile Computing*, vol. 6, pp. 943–959, 2007.

[17] C. Lam and et. al, "Fiber optic communication technologies: What's needed for datacenter network operations," *Communications Magazine, IEEE*, vol. 48, no. 7, pp. 32 –39, 2010.

[18] J. Li, C. Qiao, J. Xu, and D. Xu, "Maximizing throughput for optical burst switching networks," *IEEE/ACM Trans. Netw.*, vol. 15, pp. 1163–1176, October 2007.

[19] M. Mcgarry, M. Reisslein, and M. Maier, "Ethernet passive optical network architectures and dynamic bandwidth allocation algorithms," *Communications Surveys Tutorials, IEEE*, vol. 10, no. 3, pp. 46 –60, 2008.

[20] L. Ramakrishnan, C. Guok, K. Jackson, E. Kissel, D. M. Swany, and D. Agarwal, "On-demand overlay networks for large scientific data transfers," in *CCGrid*, 2010.

[21] W. Stallings, *Data and Computer Communications*, 8th ed. Prentice Hall, 2007.

[22] M. Welzl, *Network Congestion Control: Managing Internet Traffic*. John Wiley and Sons Ltd., 2005.

[23] M. McNierney, "College increases web speed for trial period," *The Dartmouth*, 2011. [Online]. Available: http://thedartmouth.com/2011/02/25/news/internet

[24] S. O'Malley, "Flood watch: Ohio's internet connection overflows," *http://www.ohio.edu/oit/news/ohio-internet-connection-overflows.cfm*, March 2011.

[25] J. Roettgers, "Ohio unversity blocks netflix, backpedals," *http://gigaom.com/video/ohio-university-blocks-netflix*, March 2011.