

# StorNet: Co-Scheduling of End-to-End Bandwidth Reservation on Storage and Network Systems for High-Performance Data Transfers

Junmin Gu, Dimitrios Katramatos, Xin Liu, Vijaya Natarajan, Arie Shoshani, Alex Sim, Dantong Yu, Scott Bradley, Shawn McKee

**Abstract**—Modern scientific data-intensive applications brought about the need for novel data transfer technologies and automated tools capable of effectively utilizing available raw network bandwidth and intelligently assisting scientists in replicating large volumes of data to desired locations in a timely manner. In this paper we describe the design of StorNet, an integrated end-to-end resource provisioning and management system for high performance data transfers that can operate with heterogeneous network protocols and storage systems in a federated computing environment. StorNet allocates and co-schedules storage and network resources involved in data transfers. It is based on existing Storage Resource Manager, TeraPaths, and OSCARS capabilities. StorNet provides data intensive applications with the capability of predictable, yet efficient delivery of data at rates of multiple gigabits/second, bridging end-to-end advanced storage and network technologies in a transparent way.

**Index Terms**—co-scheduling, bandwidth reservation, data transfer.

## I. INTRODUCTION

**D**ATA-INTENSIVE application communities, including high energy and nuclear physics, astrophysics, climate modeling, nanoscale materials science, and genomics, just to name a few, are expected to generate exabytes of data over the next 5 years. Such data must be transferred, analyzed, and visualized by geographically distributed teams of scientists. This expectation of explosive growth in stored data and globally distributed data processing needs, underpinned by the maturing Grid and cloud computing technologies, has generated critical requirements for new predictable and well-behaved data transfer technologies and automated tools. To expedite scientific discoveries, these data transfer tools need to intelligently assist scientists in replicating large volumes of data whenever and wherever necessary. Existing data transfer techniques face unprecedented challenges in handling not only the sheer volume of data, but also the heterogeneous environment where data are imported from and exported to. An obstacle to managing these challenges is the inability to provide end-to-end bandwidth guarantees from the source storage systems to the destination storage systems. Furthermore, technology advancements give rise to performance improvements while also increasing the complexity of resource management and provisioning. Recently, two major research and education networks, ESnet,

run by the US Department Of Energy (DOE), and Internet2, have been enhanced with advanced dynamic circuit switching technologies and network resource reservation systems to ensure on-demand bandwidth guarantees and Quality of Service (QoS). Data storage technologies have demonstrated significant improvements as well, through the use of advanced parallel file systems that enhance I/O bandwidth, and solid state disks (SSD) that can provide read/write access as much as ten times faster than hard drives. The StorNet project addresses the end-to-end resource provisioning and management issues encountered in automated data transfers by seamlessly integrating advanced network resource reservation capabilities with enhanced storage resource management (SRM) technology.

## II. MOTIVATION AND BACKGROUND

Common requirements among today's experimental science applications that are of critical importance to the mission of large experimental facilities, such as the Large Synoptic Survey Telescope (LSST) [1], the Large Hadron Collider (LHC) [2], the Spallation Neutron Source (SNS) [3], the Advanced Photon Source (APS) [4], and the Relativistic Heavy Ion Collider (RHIC) [5], are: (i) intensive data transfers; (ii) remote visualizations of datasets and on-going computations; (iii) computational monitoring and steering; and (iv) remote experimentation and control. These applications utilize a wide variety of platforms, hardware, network, storage media, and software components to deliver the critical data storage functionality: file servers (NFS and AFS), various FTP file servers, mass storage systems, relational databases, and web servers for serving files and on-line streaming video. Storage and processing of raw data takes place at geographically distributed computing facilities; sharing data across the globe is realized through transfers over high-speed networks. Because the default network behavior is to treat all data flows equally, data flows of higher priority and/or urgency may be adversely impacted by competing data flows of lower priority. In distributed data-intensive environments, this can be a major problem that significantly degrades the effective "goodput" of the overall system. The policies and priorities of user communities cannot be effectively expressed or implemented in the network, except by very labor-intensive and error-prone human intervention.

There is an evident need for coordination between storage and network systems to better service the data transfers of the

user community. From the network perspective, the capability to prioritize, protect, and regulate the various data flows becomes of high importance, because such a capability can be used for deterministically scheduling network resources to support user community priorities and, furthermore, co-schedule associated resources such as storage systems. From the storage systems perspective, the source and destination storage systems need to have adequate bandwidth and storage allocation to take advantage of the network capabilities and increase the reliability and predictability of a transfer. Furthermore, data transfers are typically long-duration and transient failures are likely to occur, so failure detection and recovery mechanisms are also necessary.

The primary goal of StorNet is to achieve the coordination of storage and network resources by taking advantage of existing systems, some already used in production, making them interoperable, and augmenting their functionality. In addition to storage resource provisioning coordination between source and target storage systems, there needs to be bandwidth provisioning coordination between the storage systems and the underlying network resources. The systems used by StorNet are LBNL's storage resource manager known as the Berkeley Storage Manager (BeStMan) [6], BNL's TeraPaths end-to-end virtual network path reservation system [7], and ESnet's On-demand Secure Circuits and Advance Reservation System (OSCARS) [8] network provisioning tool supported by both ESnet and Internet2.

### III. EXISTING TECHNOLOGIES USED IN STORNET

#### A. SRM-BeStMan

When dealing with storing large amounts of data, scientists need to interact with a variety of storage systems, each with different interfaces and security mechanisms, and to pre-allocate storage to ensure that data generation and analysis tasks can take place successfully. To accommodate this need, the concept of Storage Resource Managers (SRMs) was developed at LBNL [9,10].

SRMs are middleware components whose function is to provide a common access interface, dynamic space allocation, and file management for shared distributed storage systems. The SRM interface was standardized, and the specification led to the development of multiple SRMs that interoperate with each other by various institutions around the world [11,12,13]. SRMs are designed to provide support for storage space reservations, flexible storage policies, lifetime control of files to manage space cleanup, and performance estimation. The most recent version of an SRM developed at LBNL, is called the Berkeley Storage Manager, or BeStMan. BeStMan is designed in a modular fashion, so that it can be adapted easily to different storage systems (such as disk-based systems, mass storage systems, and parallel file systems, such as Lustre) as well as use different transfer protocols (including GSIFTP, FTP, BBFTP, HTTP, HTTPS). BeStMan is implemented in Java in order to be highly portable. It supports the basic SRM functions, and in addition, directory management and brokering service for accessing files in the distributed system. It manages queues of multiple requests to get or put files into spaces it manages, where each request can be for multiple files

or entire directories. When managing multiple files, BeStMan can take advantage of the available network bandwidth by scheduling multiple concurrent file transfers.

#### B. TeraPaths

The TeraPaths project [7] at Brookhaven National Laboratory (BNL) has been developing a host-to-host network resource reservation tool. TeraPaths utilizes a combination of DiffServ-based LAN QoS with WAN MPLS tunnels and dynamic circuits to establish host-to-host virtual paths with QoS guarantees. These virtual paths prioritize, protect, and regulate network flows in accordance with site agreements and user requests, and prevent the disruptive effects that conventional network flows can bring to one another.

Providing a host-to-host virtual network path with QoS guarantees (e.g., guaranteed bandwidth) to a specific data flow requires the timely configuration of all network devices along the route between a given source and a given destination. Typically, such a route passes through multiple administrative domains and there is no single control center able to perform the configuration of all devices involved. TeraPaths achieves this goal by directly configuring end-site LAN domains and interfacing with OSCARS for WAN domains.

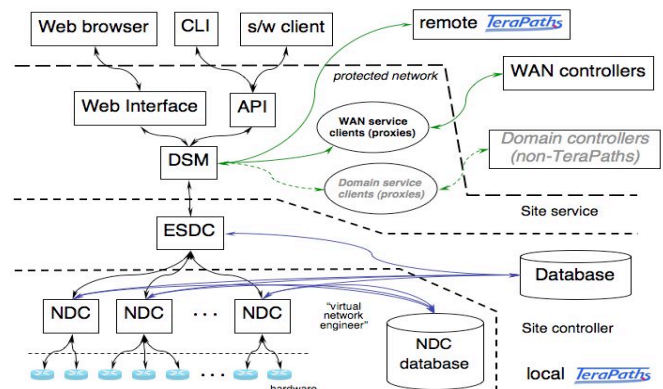


Figure 1: The TeraPaths site architecture.

The TeraPaths system [14] has a fully distributed, layered architecture (see Figure 1) and interacts with the network from the perspective of end-site users/applications. The local network of each participating end-site is under the control of an End-Site Domain Controller module (ESDC). The site's network devices are configured by one or more Network Device Controller modules (NDCs). NDCs play the role of a "virtual network engineer" in the sense that they securely expose a very specific set of device configuration commands to the ESDC module. The core of each TeraPaths site service is a Distributed Services Module (DSM). The DSM has the role of coordinating all network domains along the route between two end hosts (each host belonging to a different end-site) to timely configure the necessary network segments that establish a desired host-to-host path. The DSM interfaces with the ESDCs of its own and other remote sites to configure the path within the end-site LANs (direct control), and furthermore interfaces with WAN controlling software (OSCARS Inter-Domain Controllers - IDCs) to bring up the necessary path segments through WAN domains (indirect control).

### C. On-demand Secure Circuits and Advance Reservation System (OSCARS)

OSCARS [8] is a guaranteed bandwidth provisioning system for DOE’s ESnet standard IP network and advanced Science Data Network (SDN). It was designed specifically to meet the requirements of data-intensive scientific applications through dynamically provisioned virtual paths with guaranteed QoS, and has effectively demonstrated that an end site can reserve bandwidth within ESnet to accommodate deadline-based scheduling. OSCARS initially provided guaranteed bandwidth circuits within ESnet in the form of MPLS tunnels (layer 3). Through the collaboration between ESnet and Internet2, the system evolved into a more general Inter-Domain Controller (IDC) which provides not only MPLS tunnels within ESnet, but also guaranteed bandwidth layer 2 circuits within and between ESnet’s Science Data Network (SDN) and Internet2’s Dynamic Circuit Network (DCN) [15].

## IV. HIGH-LEVEL ARCHITECTURE

StorNet is a versatile, end-to-end, performance-guaranteed data transfer system based on an existing storage resource management system (BeStMan/SRM), and a tool for providing virtual paths with bandwidth guarantees (TeraPaths). By integrating and optimizing storage and network bandwidth provisioning and storage space reservation together in an end-to-end manner, StorNet provides data transfer applications with guaranteed and predictable QoS. At the core of the project is a flexible protocol that enables BeStMan SRMs to interoperate with TeraPaths instances, which in turn interoperate with OSCARS Inter-Domain Controllers (IDCs), and negotiate the reservations of virtual network paths with guaranteed QoS parameters spanning multiple network domains. Subsequently, SRMs use the established virtual paths to perform data transfers with increased reliability and predictability in terms of bandwidth utilization and transfer duration. The protocol also aims to provide users and applications with capabilities to detect and recover from failures, not only within the network, e.g., due to failed connections, but also within the storage sites, e.g., due to malfunctioning hardware/software.

In the StorNet framework, we follow a layered approach to compose the functionality of multiple systems and achieve the overall goal of efficient, high-performance data transfers, as shown in figure 2. The framework comprises four layers: 1) the data plane consists of disk and/or tape storage systems, site LANs and WAN backbone; 2) the control plane includes storage resource schedulers, LAN QoS provisioning and circuit utilization systems, and WAN backbone bandwidth and circuit provisioning systems based on MPLS/GMPLS traffic engineering; 3) the management plane monitors resource functionality and performance, diagnoses faults and coordinates fault recovery attempts; and 4) the service plane which reserves resources and exposes the functionality of individual systems, while also providing authentication and authorization. The service plane interacts with the control plane to dedicate data plane resources to meet data transfer and storage requirements based on application requests. In Figure 2, the horizontal direction represents end-to-end functionality. Each plane’s components serve an end-to-end

goal. The data plane is the vehicle of a data transfer between end site storage systems via the interconnecting network. The management plane provides an “enterprise” view of performance metrics, which can be used for diagnosing problems. The control plane enacts the service plane’s directions into system configurations that physically provision the required resources. The service plane negotiates the reservation of resources across domains so that an application request can be accommodated. The vertical direction represents system integration. BeStMan/SRM schedules and coordinates access to storage systems and data transfers at the behest of an application request. TeraPaths schedules end-site LAN bandwidth and configures LAN devices to dedicate this bandwidth to specific network traffic. OSCARS schedules and provisions bandwidth in the WAN domains that interconnect the end-sites.

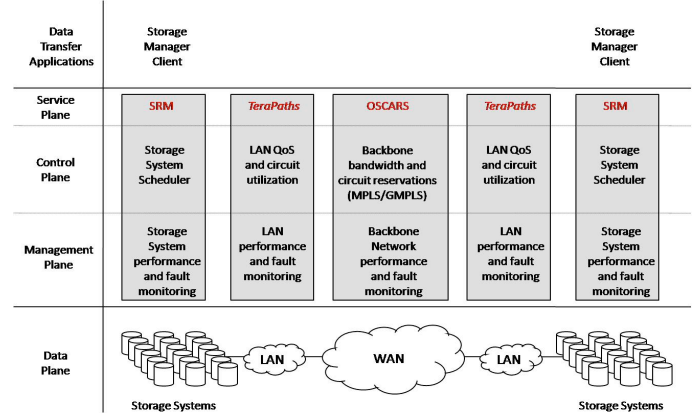


Figure 2: Efficient, high performance data transfers require interaction and cooperation among components within a set of conceptual functionality planes.

System interactions take place at the service plane layer. Triggered by a client’s request, end-site SRMs first coordinate between them to reserve storage space, and decide the parameter space (in terms of maximum bandwidth and maximum time to completion) that satisfies the request. This parameter space is then passed to TeraPaths as a request for network bandwidth reservation. TeraPaths instances coordinate between themselves to match the SRM request to LAN resource availability. Subsequently, TeraPaths generates corresponding requests for WAN bandwidth reservations and submits them to OSCARS. When multiple WAN domains are involved, OSCARS IDCs coordinate in a daisy-chain manner to establish the path interconnecting the end-sites; however, this is done transparently, i.e., TeraPaths only interacts with one IDC (see Figure 3).




Figure 3: StorNet workflow

## V. FUNCTIONALITY DESIGN

The focus of this section is on the design and enhancements of the components in the service plane, especially the communication and coordination of bandwidth between BeStMan, TeraPaths, and OSCARS.

### A. Resource Co-Scheduling

In extreme scale science environments, the resources located at each site, such as computing power and storage space, have to be allocated jointly with network resources to achieve a cost-effective and reliable data transfer and sustain the desired overall performance of distributed tasks. For instance, a site with rich storage resources may not be a good candidate for data backup if its network connectivity with other sites is poor. In such an environment where users share and compete for resources, it is critical to achieve efficient resource utilization with suitable co-scheduling schemes. StorNet addresses a general Resource Co-Scheduling (RCS) problem: given a set of limited resources of different types and a variety of requests from data-intensive applications, determine how to optimally allocate and schedule the resources required by each application. For example, consider an application performing a time-constrained end-to-end data transfer. To reliably transfer data from source disks to destination disks over the network at known rates and meet its deadline, this application may simultaneously require a bandwidth-guaranteed network circuit and a number of dedicated CPUs and hard disks. We therefore need to jointly allocate and co-schedule all required types of resources.




Figure 4: Bandwidth Availability Graphs

In the context of the StorNet project, we have developed for this purpose an analytical model of resource co-scheduling, based on the concept of an end-to-end Bandwidth Availability Graph (BAG). We assume that the utilization of each resource type can be scheduled by advance reservations with specific start and end time and constant bandwidth allocation for their duration. The bandwidth allocation of such a set of reservations can be aggregated and subsequently subtracted from the maximum bandwidth availability for the overall time period to yield the BAG for the resource of interest (see Figure 4a). The maximum availability can vary with time, but typically can be considered constant, at least within known

time intervals. As such, a BAG is a step function. For a storage system, for example, the maximum availability could be the total achievable transfer rate, and for a network domain the maximum achievable bandwidth. Individual BAGs can be intersected to express the minimum availability of the initial BAGs at any given time, which provides the overall availability of resources across any number of systems (see Figure 4b). The intersection of all BAGs of source and destination storage systems and interconnecting network domains yields the end-to-end BAG.

Subsequently, a new request for reserving that resource can be represented by a rectangle (see Figure 4c). If the rectangle fits into the overall BAG, then the request can be satisfied. A request may be flexible in terms of start time, duration, and/or bandwidth so that the rectangle can be modified to fit into the graph (see Figure 4d). In the latter case, the area of the rectangle represents the total volume of data to be transferred, and any modification to the start time, duration, and/or bandwidth must result in a rectangle with the same area as the initial one. The objective of fitting the request rectangle is to obtain a solution (i.e., a set of reservation parameters acceptable across all systems) that optimally satisfies the request. As optimal, we define a solution that satisfies the request according to the requestor's preferences. We have primarily considered the cases of shortest transfer duration and earliest finish time.




Figure 5: Finding a solution

Fitting the request rectangle can be approached as a variation of the problem of finding the largest rectangle under a histogram with  $n$  adjacent rectangles, which can be solved in  $O(n)$  time [16]. More specifically, a BAG is represented by a sequence of  $n$  windows  $[s_i, e_i]$ , where  $s$  is the start time and  $e$  the end time, each with constant bandwidth  $b_i$ . Firstly, we observe that an optimal solution can always start from some  $s_i$ . If an optimal solution starts elsewhere within a window, we can move the whole solution (rectangle) to the left until it starts from an  $s_i$  point; the solution will have identical duration but earlier finish time (see Figure 5a). Secondly, we can obtain exactly  $n$  rectangles  $\{start[i], end[i]\}$  with bandwidth  $b_i$ , where  $start$  and  $end$  denote the start and end time of the largest rectangle containing window  $i$  (see Figure 5b). We can then fit a request with a given same data volume in these  $n$  “largest” rectangles and choose the one with the shortest duration or the earliest finish time, depending on the request's preferences.

Based on the above observations, we can obtain the optimal solution in  $O(n)$  time using one of the known algorithms for the problem of finding the largest rectangle under a histogram [17]. In other words, given a BAG, it will take  $O(n)$  time to obtain  $start[i]$  and  $end[i]$  of the largest rectangle corresponding to every window in the BAG. An example of such an algorithm that uses a stack is shown in figure 6. The idea is that the stack holds a series of windows

that have increasing heights; windows are enumerated from left to right.

```

Append a dummy window with zero height and zero width to the list of windows
Merge successive windows of equal bandwidth (enforce  $b_q < b_p, q=p+1, p=1..n$ )
For each window  $i$  from left to right {
  If (stack.empty() == true or bandwidth of current window  $b_i >$  stack.top().bandwidth) {
    push window  $i$ , set start[ $i$ ] =  $s_i$ 
  } else {
    repeat {
      pop window  $j$ , set end[ $j$ ] =  $s_j$ , set  $k = j$ 
    } until (stack.empty() == true or  $b_i \geq$  stack.top().bandwidth)
    if (stack.empty() == false and stack.top().bandwidth ==  $b_i$ ) {
      set start[ $i$ ] = start[stack.top()]
    } else {
      set start[ $i$ ] = start[ $k$ ], where  $k$  is the last popped window
    }
    push window  $i$ 
  }
}

```

Figure 6: Pseudo-code for stack-based largest rectangle algorithm

StorNet approaches schedule negotiation in a top-down direction across systems, i.e., narrowing down the solution space is first performed at the SRM level, then at the TeraPaths level, and finally at the OSCARS level. This is done for two major reasons: firstly, because the availability of resources within each system must take into account the aspects of system-wide policies and user privileges; and secondly, because the amount of effort for figuring out a solution is reduced. Although BAG intersection is commutative, using a separate scheduling component is not feasible, as it would require systems belonging to different administrative domains to reveal non-public information in a bottom-up manner. For example, it cannot be expected from OSCARS to reveal to TeraPaths all schedule information pertaining to a network path of interest. In contrast, in the top-down direction one system passes to another only non-sensitive information required to obtain solutions satisfying the original request. An additional incentive for minimizing the candidate solution set is that the current implementation of OSCARS does not support negotiation with BAGs and candidate solutions have to be tried one by one in a time-expensive trial-and-error manner. Reducing the number of options based on previous constraints reduces the search space and therefore the interaction with OSCARS.

### B. BeStMan Functionality Enhancements

The data transfer protocols currently used by BeStMan, such as GridFTP, assume best-effort IP networks, and improve performance with a large number of TCP streams for long, round-trip connections. Fairness and efficiency are adversely affected by such a brute force data transfer method. The primary goal of StorNet is to provide data transfers with QoS guarantees and to move away from the best-effort data transfer paradigm that does not provide delivery time assurance. To support network and storage co-scheduling, the existing data transfer module in BeStMan is being extended to reserve end-to-end network bandwidth and intelligently optimize storage space and network bandwidth allocation, thus increasing transfer reliability. This extension will reduce the “impedance mismatch” between end user data transfer applications, storage, and the network. In order to keep track of bandwidth reservations and commitments, BeStMan is also being enhanced with a backend database service. This will provide

persistent store for tracking user requests, storage space allocations, and bandwidth allocations.

The enhanced BeStMan is designed to achieve the best solution for user requests. Users can specify whether they prefer earlier time solutions or shortest transfer, and they provide BeStMan with a desired time of completion. The BeStMan at the target site (pulling the data) also needs to have the logic to communicate with the source BeStMan to find out what is its bandwidth availability. The source BeStMan returns the availability “graph” for the requested period of interest (i.e. till maximum time), in the form of a sequence of windows. The target BeStMan then finds a common schedule, and provides that to TeraPaths. The API for BeStMan-TeraPaths interaction is described in section V.D.

### C. TeraPaths Functionality Enhancements

To accommodate the functionality required for StorNet, TeraPaths is being enhanced along two main directions: interaction with BeStMan and core extensions to support negotiation between end-sites and with OSCARS. Communication and coordination with BeStMan is supported by a BeStMan-to-TeraPaths (StorNet) API module that interprets and validates BeStMan requests and passes them along to the main system through the TeraPaths API. The former API is essentially a wrapper of the latter. The choice of using an API wrapper allows us to standardize and simplify the interaction between BeStMan and TeraPaths, so that future revision of one API will not necessarily affect the other API or BeStMan’s client. Core extensions to TeraPaths are primarily necessary for supporting negotiation between end-site instances through BAGs, calculation of solutions spaces by fitting requests into intersected BAGs, and negotiation with OSCARS by applying a trial-and-error approach on the set of candidate solutions obtained from the fitting process.

### D. BeStMan-TeraPaths Web-Services Interface Design

We designed a BeStMan-TeraPaths web-service interface to describe the functions necessary for BeStMan server to request network bandwidth from TeraPaths service. The goal of the API is to enable BeStMan to negotiate bandwidth with TeraPaths. The important functionalities reflected in the interface are bandwidth reservation, commitment, modification, and cancelation. The interface also includes status check and time-out extension. Necessary information, such as data volume, source and target resource availability, resource time frames, and other attributes, is provided to TeraPaths when requesting network bandwidth. In Figure 3, we showed the sequence of communications between the components. In Step (1), which involves getting the request from the application, a list of files or a directory are provided as well as source and destination. In addition, a window of desired start time and maximum completion time is provided. In step (2), the source and target BeStMan servers communicate with each other to reserve storage space and to determine the maximum bandwidth they can both use during the requested window. Once this is determined, BeStMan communicates with TeraPaths in step (3), and provides a sequence (start time, end time) of non-overlapping windows, and maximum bandwidth for each. Based on this information,

TeraPaths schedules its local area resources, and then negotiates with OSCARS for wide area resources. The communications in steps (4) and (5) are internal to TeraPaths, and only the resulting reserved window is communicated to BeStMan. A typical scenario is that BeStMan first tries to make a temporary network bandwidth reservation. If such a reservation is possible, TeraPaths returns a request token, along with an expiration time and available windows for the available resources. Once BeStMan determines that it can work with the result from TeraPaths, it commits the reservation to lock in the network resources. Otherwise, BeStMan modifies its input and submits a new request.

## VI. EARLY RESULTS

The StorNet functionality was demonstrated at SuperComputing 2010 using an early prototype implementation of the architecture. The demonstration included actual 10 GB file transfers between BNL and University of Michigan using the current TeraPaths testbed (see figure 7). Heavy interference traffic allowed best effort transfers to only reach roughly 8 MB/s (7c), while transfers with StorNet could be tuned to desired levels of bandwidth (7a,b,d) unaffected of network congestion. Because of the low performance of best effort transfers, a smaller file of 1 GB size was used to save demo time.

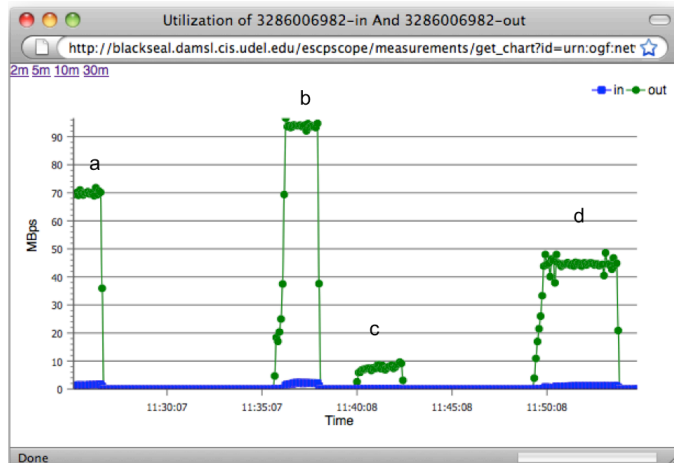


Figure 7: StorNet SC'10 demo: network interface traffic at sending host: (a), (b), (d) transfer of a 10 GB file with 70MB/s (partially displayed), 95 MB/s, 45 MB/s reservations respectively; (c) transfer of a 1 GB file w/o reservation.

## VII. SUMMARY AND FUTURE WORK

Effective and robust data transfer is essential to current day scientific applications, and is a major concern for future scientific work as the volumes of data collected and shared grow exponentially. In order to address this problem, two aspects need to be supported: 1) ways to reserve and guarantee bandwidth in network and storage, and 2) ways to coordinate and synchronize bandwidth reservations in all components from source to destination. Thus, storage systems need to be instrumented to support such bandwidth reservations, and coordinate with local area and wide area network bandwidth provisioning. For this purpose, there need to exist components that can control the bandwidth reservations, provision the bandwidth, and ensure that the allocated bandwidth is used effectively. In this work, we take advantage of existing storage

and local network middleware technologies (called BeStMan<sup>1</sup> and TeraPaths, respectively) to pursue this goal. Bandwidth provisioning in the WAN is realized by having TeraPaths negotiate with ESnet's OSCARS provisioning system. This coordinated approach is achieved by enhancing the existing middleware systems with APIs for negotiating end-to-end bandwidth reservations and obtaining monitoring information. The design of such APIs has been described.

In the current design we are targeting the reservation of a single transfer window per request, i.e., of a window with constant bandwidth (height) and specific time duration (width). Given a busy network, there can always be a possible solution if such a window is scheduled into the future beyond current commitments. A more attractive solution is, of course, to allow for multiple windows per request, filling available reservation gaps. This problem is much more challenging to optimize and manage. We allowed for such a possibility by having multiple windows per request in the APIs, and plan to address this possibility in the future, after the single window strategy is fully implemented.

## ACKNOWLEDGMENT

This work is supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 (to LBNL) and DE-AC02-98CH10886 (to BNL).

## REFERENCES

- [1] LSST: <http://www.lsst.org/lsst>
- [2] LHC – The Large Hadron Collider: <http://lhc.web.cern.ch/lhc/>
- [3] Spallation Neutron Source (SNS): <http://neutrons.ornl.gov/aboutsns/aboutsns.shtml>
- [4] Advanced Photon Source (APS): <http://www.aps.anl.gov/>
- [5] Relativistic Heavy Ion Collider (RHIC): <http://www.bnl.gov/rhic>
- [6] Berkeley Storage Manager: <http://sdm.lbl.gov/bestman/>
- [7] D. Yu and D. Katramatos, TeraPaths: <http://www.terapaths.org>
- [8] OSCARS: <https://oscars.es.net/OSCARS/docs/>
- [9] Shoshani, A., Sim, A., and Gu, J.: Storage Resource Managers: Middleware Components for Grid Storage. In: Nineteenth IEEE Symposium on Mass Storage Systems, 2002
- [10] Shoshani, A., Sim, A., and Gu, J.: Storage Resource Managers: Essential Components for the Grid. Chapter in book: Grid Resource Management: State of the Art and Future Trends, Edited by J. Nabrzyski, J. M. Schopf, and J. Weglarz, Kluwer Academic Publishers, 2003
- [11] A. Sim, A. Shoshani, F. Donno, J. Jensen: Storage Resource Manager Interface Specification V2.2 Implementations Experience Report. GFD.154, Open Grid Forum, Aug. 2009
- [12] A. Sim, A. Shoshani (Editors): The Storage Resource Manager Interface Specification Version 2.2. GFD.129, Open Grid Forum, Document in Full Recommendation, Feb. 2008
- [13] Lana Abadie, Paolo Badino, Jean-Philippe Baud, Arie Shoshani, Alex Sim, et al.: Storage Resource Manager version 2.2: design, implementation, and testing experience. In: Conference for Computing in High Energy and Nuclear Physics, 2007
- [14] Dimitrios Katramatos, Bruce Gibbard, Dantong Yu, Shawn McKee: TeraPaths: End-to-End Network Path QoS Configuration Using Cross-Domain Reservation Negotiation. In: 3rd International Conference on Broadband Communications, Networks, and Systems (BROADNETS 2006), 2006
- [15] Internet2 Dynamic Circuit Network (DCN): <http://www.internet2.edu/network/dc/>
- [16] Carroll Morgan: Chapter 21: The Largest Rectangle under a Histogram. Programming from Specifications, 2<sup>nd</sup> edition, Prentice Hall International (UK) Limited, October 1998
- [17] ACM Collegiate Programming Contest 2003/2004: <http://www.informatik.uni-ulm.de/acm/Locals/2003/html/judge.html>